


Cyber-Catastrophe Insurance-Linked Securities On Smart Ledgers



November 2018



Cyber-Catastrophe Insurance-Linked Securities On Smart Ledgers

November 2018

Sam Carter

Carter Research Ltd.

Professor Michael Mainelli

Executive Chairman, Z/Yen Group

“Are you interested in Catastrophism?” asked the wondering Yankee.

“I’m interested in catastrophes; and there are going to be some,” replied his companion gloomily. “Mine’s a filthy trade, and I never pretend it isn’t.”

G.K. Chesterton, “The Strange Crime of John Boulnois”.

Foreword

Insurance has evolved over centuries. Little by little, practitioners have broadened their understanding of what risk is, how it is distributed, and how to cover it. This increase in insight is often driven by innovations in technology and commerce. As new risks emerge, it is up to insurers to apply their old strategies in new ways, or think up novel strategies. Cyber risk, the risk to people and businesses posed by information & computing technology (ICT) is a multi-faceted risk, potentially leading to loss of data, revenues, direct physical harm, or reputation. Many areas of cyber risk are hard to insure because third-party or consequential damages are hard to estimate.

One area that could be covered is 'business interruption'. Business interruption is an existing class of insurance exacerbated by cyber attacks or outages. Business interruption could be extended more easily to cyber interruptions if there was an ability to reinsure such risk. However, reinsurance in turn requires an ability to show that the risk of a cyber catastrophe, the risk that a significant portion of ICT is out of action, is covered for a reinsurer. Such catastrophe risks are more and more addressed by Insurance-Linked Securities (ILSs). The outsourcing of catastrophe risk to the financial markets, through ILSs, is a welcome addition to the insurer's arsenal of risk reduction measures. It is a natural extension of historical trends in the insurance industry.

The world of computers surrounds us like an ecosystem. Each machine in the system becomes more and more complicated every year, and the networks which link the machines become more tangled. The interdependence of the machines on each other therefore becomes more and more difficult to understand. As the complexity of our computer networks grows, they become more sensitive to fluctuations in initial conditions. Catastrophic computer and network failure becomes a real risk. Yet, catastrophic failure might be easily expressed as a percentage of specified computers unavailable for a specific length of time. Perhaps 25% of computers for the UK's critical national infrastructure unavailable for 4 hours; or 40% of the Standard & Poor's 500 websites out of action for more than 2 days.

This report examines the history of insurance and re-insurance, looking at how technological and social changes have spurred the evolution of thinking about risk. It looks at the evolution of Insurance-Linked Securities and Catastrophe Bonds, then discusses Cyber risk. Within the rather broad definition of 'cyber',

the report unpicks those elements that might be appropriate for a market-driven solution. ILSs might well be one solution for mitigating the risk of cyber-catastrophe, so what is the current market like, and what might the structure of such securities be.

Finally, the report explores a novel technical solution, using a dynamic index based on a Smart Ledger. The Smart Ledger approach allows participants to measure the current levels of network risk, and potentially trigger insurance payments. The benefit of this technology is that it provides high security and large degrees of flexibility in structure, yet prevents the rise of an over-weening central third party, thus encouraging competition and innovation in the provision of cyber insurance direct to clients.

I welcome this report for taking cyber catastrophe and cyber business interruption seriously and demonstrating one way in which the insurance industry can innovate, and commend this report to practitioners, regulators, and policy makers to stimulate further thinking.



Julian Enoizi

Chief Executive, Pool Reinsurance Company Limited

Contents

Foreword	3
Executive Summary.....	7
1. The Evolution of Insurance.....	10
A. What Is Insurance?.....	10
B. A Brief History Of Data	11
C. Risk Factors	13
D. New Markets For Insurance	13
E. The Birth Of Reinsurance	15
F. Terrorism And Pool Re	17
2. Insurance-Linked Securities & Catastrophe Bonds.....	20
A. History & Motivation.....	20
B. Structure	21
C. Triggering And Parametrisation	23
D. The ILS Market	26
E. London’s Changing Role	28
F. Terror-Linked ILS.....	28
3. Cyber Risk	30
A. Definition	30
B. Current Insurance Offerings.....	32
C. Attacks Causing Business Interruption	34
D. Current Reinsurance Capacity	39
E. Current Cyber Insurance Technology	40
4. Cyber-Catastrophe And ILS	42
A. A Gap In The Market	42
B. More Data More Models.....	43
C. Structuring The Risk	44
D. Triggers	45
E. New Markets For ILS	46
F. The Future Of ILS Cyber.....	47
5. A Smart Ledger Solution.....	50

A.	The Network Availability Index.....	50
B.	Polling The Network.....	51
C.	Constructing The Index	54
D.	Constituent Weighting	55
E.	Proposed Architecture	56
F.	A Broadcast/Receive Solution	59
G.	The Benefits Of This Approach	62
6.	Further Applications.....	64
A.	Polling	64
B.	Publishing Indices.....	66
C.	Triggers	69
	Conclusion	70
	Principal Authors	71
	Acknowledgments	73

Executive Summary

In life, as in commerce, when faced with risks, we have four possible courses of action available to us:

- Avoid
- Reduce
- Transfer
- Retain or Accept

If at all feasible, we would like to “Avoid” or “Reduce” the risks. There are times, however, when this is either too expensive, impractical, or simply impossible without a crystal ball. At such times, the insurance industry can be called upon to help us with a “Transfer” strategy.

The theory of insurance has evolved over centuries. Little by little, practitioners have broadened their understanding of what risk is, how it is distributed, and how to manage it. Innovations in technology and commerce drive these new insights. As new risks become apparent, it is up to insurers to work out how to apply their existing strategies in new ways, or to think up entirely new strategies.

The nineteenth century saw massive growth of population and the development of heavy industry brought into focus the concept of catastrophe risk - where huge numbers of people are affected by a single event. Catastrophe risk spawned the reinsurance industry, and with it, novel ideas about quantification, spreading and transfer of risk.

Nowadays, insurers face new challenges. The world of computers surrounds us like an ecosystem. Each machine in the system becomes more and more complicated every year, and the networks which link the machines become more tangled. The interdependence of the machines on each other therefore becomes more and more difficult to understand.

Since computers are logical constructions, one would have thought that the causes of their failures would be easy to track and understand in a step-by-step fashion. In fact, this turns out not to be the case. The emergent complexity of a networked system is far greater than the sum of its parts.

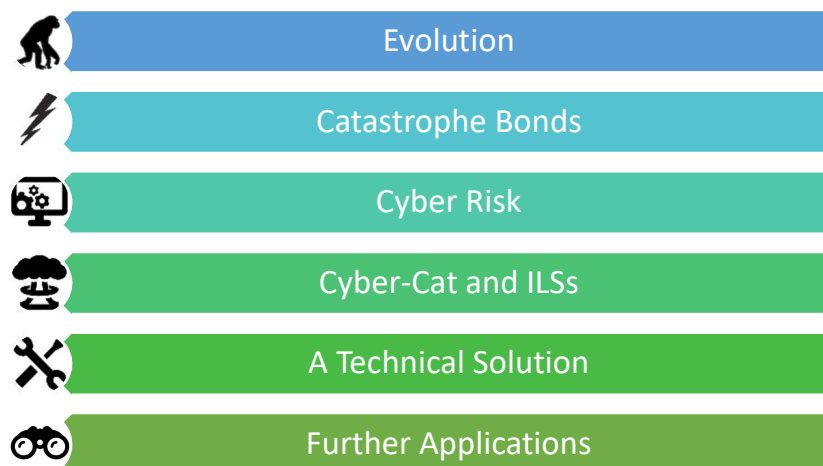
One characteristic of complex systems is their potentially chaotic nature. We are familiar with the notion that a butterfly flapping its wings in the Maldives

can cause a hurricane in the Atlantic. This sensitivity to initial conditions makes the weather very hard to predict even a couple of days ahead. Similarly, as the complexity of our computer networks grows, they become more sensitive to fluctuations in initial conditions. Catastrophic failure becomes a real risk.

Is the insurance industry ready for a cyber-catastrophe? What strategies do we need to put in place right now in order to make sure that, when the event comes to pass, help is available to the victims?

In this report, we will argue that the outsourcing of catastrophe risk to the financial markets, through the use of Insurance-Linked Securities, is a welcome addition to the insurer's arsenal of risk reduction measures, and that it is a natural extension of historical trends in the insurance industry.

The report is divided into six chapters:



First, we examine the history of insurance, starting with medieval maritime and early life insurance, up to the introduction of motor and liability insurance. We look at how the re-insurance industry was born in the ashes of great city fires, and we conclude by discussing how reinsurers are attempting to manage the risks associated with acts of terrorism.

In the second chapter, we look at the current market for Insurance-Linked Securities and Catastrophe Bonds, and how they have evolved over time since their introduction in the early 1990's.

Then we examine cyber risk. What is it and how is it tackled? We look at the extremely broad definition of “cyber”, and attempt to unpick which elements in the sector might be appropriate for a market-driven solution.

We then ask whether the Insurance-Linked Security (ILS) is a good solution for mitigating the risk of cyber-catastrophe. What is the current ILS market like, and what might the structure of such securities be?

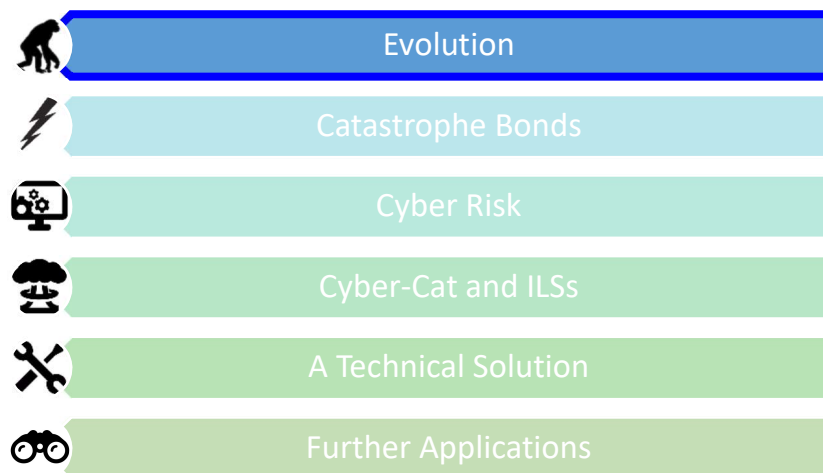
With all of the above in mind, we then explore a possible technical solution. We propose a novel dynamic index which will allow participants to measure the current levels of network risk, and potentially trigger insurance payments. Smart Ledgers - distributed databases with a super audit trail are explored in this respect. This new technology makes it possible to implement intelligent, resilient, automatic processes, ideally suited to the ILS setup.

Finally, we examine the potential commercial applications of the setup, and how it might be applied to other markets.

1. The Evolution of Insurance

Over the last few centuries, humans have come up with a variety of ways to minimise the risks inherent in all activities, particularly large commercial enterprises. The limited company, the use of collateral, the derivative security – all these are an attempt to separate the cost of risk from the main cost of the activity itself.

In this chapter we discuss how we have gone from trying to account for risk *within* industrial and maritime contracts, to being able to buy and sell coverage as an independent product. In order to be able to deal with new types of risk, we need to see how new risks have been understood and dealt with in the past.



A. What Is Insurance?

An insurance policy is a contract where an unknown future loss is exchanged for a known current premium. This financial magic works because, through the insurer, a pool of customers creates a common pool of money, allowing the community in general to share their risks with each other.

For this setup to operate with any longevity, it is vital that the pool of money is large enough, and replenished often enough, to cover all claims which could reasonably occur over a given time-frame.

In order to write a policy covering a certain event over a certain period of time, the problem is therefore to be able to specify the following numbers:

- The premium to charge. When an insured event occurs, if the insurer has charged premiums which are too low, there will be insufficient money in the pool.
- The maximum payout. The insurer must avoid paying out an amount of money which is too high. For this reason, an insurance policy will usually specify a maximum amount which the insurer is willing to pay, thus limiting the loss to the pool.
- The minimum payout. The insurer also limits the minimum amount paid, to avoid having to payout many small sums of money. The reasoning is that smaller losses occur often, and having to attend to them all would drain the pool quickly – and vastly increase administrative costs.

The problem of calculating the relationship between these values is one which has interested mathematicians, statisticians and financiers since the beginning of the industry.

B. A Brief History Of Data¹

Insurance was developed as a by-product of the funding of sea voyages. There is evidence that in the ancient world, when merchant voyages were financed, the agreements acknowledged the risks of storms or piracy, and were structured with guarantees and interest payments to cover the risk.

The first contract we recognise as a modern insurance policy dates from 1350. A shipment of wheat from Sicily to Tunis was insured for 300 florins at a premium of 18%, for which the insurer, Leonardo Cattaneo, undertook to assume all risks. This nascent insurance market was centred on Florence for some centuries, but growing interest in London caused the Privy Council to establish the Office of Assurances at the Royal Exchange in 1575.

An example of one of the first policies issued in London was a marine insurance policy of 1580. One Dominicke Butcher insured his 45-ton ship the *Carousse*, with all its cargo, including 903 hides and 9797 pounds of lead. The premium charged for the insurance was just 3 per cent.

¹ For more, see Lewin's excellent 2007 lecture, the slides and notes for which are to be found at: <http://www.actuaries.org.uk/research-and-resources/documents/overview-actuarial-history-slides-notes>

At this distance it is not at all clear how that premium was arrived at. The medieval understanding of probability and risk were vague. Money focuses the mind, however, and it appears that most ancient cultures had an appreciation of the relationship between the amount worth risking in some situation (be it marine funding or a dice game) and the probability of loss. It is unknown exactly how insurers and granters of annuities reasoned when setting out their premiums and prices.² One must assume that experience of the industry, and news and gossip of seasonal risk and piracy played important roles, not to mention the highly competitive insurance market which soon sprung up. The coffee shops of London (most famously Lloyd's) provided the ideal meeting spot for the swapping of deals and data.

In 1657, Christian Huygens published a work on probability which laid down some systematic formulas for the calculation of risk. This allowed for the development of a model-driven market over the subsequent centuries.

Life Insurance

The earliest life insurance policy we know of was issued in 1583 on the life of William Gibbons of London, for a term of 12 months at a premium of 8 per cent. At the time life policies could only be issued by the Royal Exchange for a maximum sum of £1,000, and up to one year at a time. In the rest of Europe the market was banned, for fear of encouraging murder.³

During the plague, life insurance was placed on a statistical footing. Large cities published "Bills of Mortality" once a week, and using this data, John Graunt produced the first life table in 1662. This was based more on his own model than on the data, and was very inaccurate.

However, the format of the table caught people's attention, and eventually, in 1693, the mathematician Edmund Halley published a more accurate version. He immediately applied his calculations to the problem of life insurance, estimating the present value of a payment in a future year which depended upon survival to that year.

² See James Franklin's 2001 study: *The Science of Conjecture: evidence and probability before Pascal*.

³ William Gibbons died nearly a year after the policy was effected, but the underwriters refused to pay, on the grounds that Gibbons *had* lived for 12 months, if a month was defined as 28 days. The Courts disagreed.

With this understanding, the insurance industry was placed on a firm statistical footing.

C. Risk Factors

The “pure premium” for an insured individual is the amount which will be required by the insurer to pay out the likely losses, not including any profit margin or costs. To calculate the pure premium, the insurer needs to consider two factors:

- Claim Probability. This is related to the number of claims expected to occur in a given year for a given individual.
- Claim Size. This is the cost associated with each claim.

Estimating these two elements is the fundamental challenge inherent in the provision of insurance. Various methods have been used to attack the problem, which have been described exhaustively in the literature.⁴

Crucially, the insurer does not consider each customer in isolation. A set of customers constitutes a pool of risk, and like any investor, the insurer needs to ensure that its investments are sufficiently diversified. To take a simple example, if an insurer’s only clients all live in close proximity to one another, the insurer would be unwise to insure them all against fire, as the risk cannot be spread so that it is borne by other customers.

This has been long known in the industry. In 1747 Corbyn Morris demonstrated mathematically that insurers should spread their wealth between a large number of risks in order to lessen the probability of going bust.

We will return to this point later when discussing catastrophe risk.

D. New Markets For Insurance

By the eighteenth century, it was clear that the pricing of insurance was wedded to the new science of probability. No insurer would write a policy without being fully cognisant of the risks involved.

⁴ An overview can be found in de Jong & Heller, 2008. *Generalized Linear Models for Insurance Data*, Cambridge Books, Cambridge University Press.

In order to estimate the probability of a future claim, it is necessary to have access to past data. In the case of life insurance, the event insured against is a certainty, but the *time* until the event is in doubt. The life insurance and annuity markets were kickstarted by Halley's drawing up of life tables. In non-life insurance, both the likelihood of the claim and its expected size need to be known in order to write a policy, and this must be based on past data.

This means that when cultural or technological innovation results in a market for a new type of insurance, a Catch-22 arises. Without data, the insurer cannot offer insurance, and very often, without insurance (or an unrealistically large amount of reserve capital) an industry cannot get going – and therefore no data can be generated.

One such market which managed to bootstrap itself into existence is the insurance of commercial satellites. For much of the history of the enterprise, it has been prohibitively expensive and highly risky. The cost of launching a satellite can be over \$400 million,⁵ and as many as one in ten satellites are either destroyed on launch, or fail within a year. This combination of factors has resulted in a multi-billion dollar market for insurers. Typically, a single satellite launch involves a large number of participating insurers to spread potential losses.

But how did this market get going? Since the launch of Sputnik in 1957 until the early 1980's, the operation of satellites was undertaken only by government agencies, who self-insured. When the commercial satellite market took off in the 1980's, the costs meant that commercial insurance was vital.

However, by this stage, the insurers had over twenty years of government-funded data to go on. Actuaries could therefore reliably tabulate frequency of failure by satellite type and launch method, as well as constructing "life" tables for satellites already in orbit. Pricing insurance therefore became possible, and the market came into existence relatively painlessly.⁶

New insurance markets can also come into being as a result of government legislation, either as an unexpected side-effect of new law, or as an explicit mandatory requirement.

⁵ According to Elon Musk, SpaceX may bring this down to as little as \$65 million.
<https://futurism.com/elon-musk-launching-a-satellite-with-spacex-is-300-million-cheaper/>

⁶ See <https://www.casact.org/pubs/forum/00fforum/00ff047.pdf> for an overview.

The liability insurance market, for example, arose as a result of the passage of Employers' Liability laws, in Germany in 1871, and subsequently in Britain in 1880. These acts ensured that injured workers would receive compensation, without having to sue their employers for negligence. The new laws, while not explicitly mandating liability insurance policies, ensured their creation.⁷

These early policies soon led to the creation of liability insurance of all kinds, from contractors' to landlords' to physicians', many of which became mandatory by law. In 1927 the state of Massachusetts legislated for mandatory liability insurance for automobiles,⁸ followed shortly by the UK in 1930.⁹

We see, therefore, that new insurance markets can get going, but without the government either shouldering a portion of the risk in the early days, as with satellites, or simply mandating that the insurance is required, as with automobiles, it is difficult for the market to bootstrap itself into existence.

E. The Birth Of Reinsurance

Reinsurance is insurance for insurance companies, and is as old as insurance itself. The earliest known agreement is from 1370, when the direct insurer for a cargo ship transferred the risk for the more dangerous part of the voyage to another insurer.

Reinsurance as we now know it arose in the 19th century as a result of a change in the fire insurance business. Before the Industrial Revolution, the provision of fire insurance was dominated by mutuals, who had the ability to assess each of their members, and also to appeal to them in the event of catastrophes. Subsequently, more and more stock companies entered the insurance field, and discovered that offering coverage in large amounts was financially hazardous, especially in areas of concentrated fire risk. These companies took to transferring some of their risk to other insurance companies, often competitors, effectively offering a kind of co-insurance.

⁷ See <https://www.irmi.com/articles/expert-commentary/how-umbrella-policies-started-part-1-early-liability-coverage>

⁸ <https://www.dmv.org/articles/history-of-car-insurance/>

⁹ <https://www.legislation.gov.uk/ukpga/Geo5/20-21/43/introduction>

In 1842, the Great Fire of Hamburg destroyed about a quarter of the city, and left 20,000 people homeless. Many local insurance companies went bankrupt. In the wake of this disaster, the first independent professional reinsurance company, the Cologne Re was set up in 1846. A similar fire in Glarus, Switzerland in 1861 led to the foundation of Swiss Re.¹⁰

The problem that the insurers were dealing with is known as “catastrophe accumulation”, where a natural disaster results in concentrated losses exceeding total premiums.¹¹

In order for the insurance industry to adequately cover catastrophic losses such as this, it was necessary to develop catastrophe models. In the early days this was done through map-based models, where fire and lightning events were represented by pins in a map. Any clustering of loss events could therefore be seen, allowing insurers to spread their risk appropriately.¹²

In the twentieth century, the modelling of catastrophe risk advanced through the development of more advanced instruments (the seismograph, the anemometer), the accumulation of data, and the availability of powerful computers to run the models. Firms like AIR Worldwide and Risk Management Solutions (RMS) emerged, specialising in modelling catastrophe risk.

However, it was not until 1989 that it became starkly apparent to the insurance industry that catastrophe modelling was vital. In September of that year Hurricane Hugo devastated parts of South Carolina, causing insured loss of \$4 billion. Less than a month later, the Loma Prieta Earthquake hit San Francisco, causing property damage estimated at \$6 billion.¹³

If insurers had not been given a shock in 1989, they were in August 1992, when Hurricane Andrew hit Southern Florida. Within hours, AIR Worldwide issued a loss estimate of \$13 billion. The final amount turned out to be \$2.5 billion more, and eleven U.S. insurers went bust.

¹⁰ <http://www.soa.org/library/newsletters/reinsurance-section-news/2009/february/rsn-2009-iss65.pdf>

¹¹ <https://www.investopedia.com/terms/c/catastrophe-accumulation.asp>

¹² https://www.springer.com/cda/content/document/cda_downloadaddocument/9780387230825-c2.pdf?SGWID=0-0-45-169675-p35030036

¹³ Seismicity of the United States, 1568-1989, Stover and Coffman, 1993.

The modelling and estimation of catastrophe risk became a growth industry, funded by both public and private money. To deal with the risk, insurers use specially designed catastrophe insurance as well as reinsurance and retrocession (reinsurance of reinsurance).¹⁴

The growth of the industry led to the establishment of various trading and data-sharing platforms. For example, in 1994, Catex was formed as a reinsurance risk exchange.¹⁵

F. Terrorism And Pool Re

More recently, insurers have been obliged to widen their definition of catastrophe, to include not only natural but also man-made disasters. In 1992, the Baltic Exchange in London was destroyed in a terrorist bombing by the Provisional IRA. As a result, the insurance and reinsurance industry refused to continue to provide terrorism cover. Acts of terrorism therefore joined acts of war as explicit exclusions in most standard insurance policies.

The UK's commercial property was therefore exposed to the risk of terrorism. It was immediately apparent that this was an unsustainable situation, potentially with serious implications for the economy. Following discussions between insurance industry and the government, Pool Re was formed, and began operations in 1993.

Pool Re is an industry mutual which allows for the pooling of risk. In their own words, their primary role is:

“...to enable the commercial market to underwrite the risk of damage to commercial property caused by an act of terrorism at relatively risk-reflective rates by mitigating their exposure to the catastrophic losses associated with major attacks.¹⁶”

As with all new insurance markets, the lack of data and the risk of large payouts meant that the market required a government backstop to get going. While it is possible to model the extent of terror damage by looking at attack types and blast radii, what is harder to predict is the frequency of attacks. Without this

¹⁴ <https://www.investopedia.com/terms/c/catastrophe-insurance.asp>

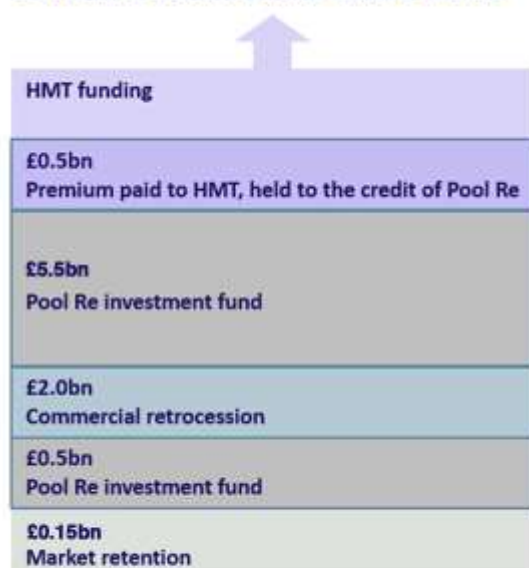
¹⁵ <http://www.catex.com/aboutus.aspx>

¹⁶ <https://www.poolre.co.uk>

data, the government must support the scheme if the industry is to avoid having to carry exorbitant amounts of reserve. The only way that Pool Re can function is with that government guarantee. As a result, Pool Re effectively have a monopoly on terrorism reinsurance. They are therefore limited to this sector and are not allowed to branch out.

After an event, the government has 21 days to declare the event an act of terrorism.¹⁷ If this is done, Pool Re may pay out to a member. Their membership comprises most of the major property insurers in London. Should a member's client experience a loss as a result of a terrorist attack, the member is only responsible for paying losses up to certain threshold. Above that, the insurer can call upon Pool Re's £6 billion reserves. If the reserves are not enough, Pool Re is entitled to call on the government, who will lend the money.

Scheme Resilience June 2016



Note: This chart excludes additional premium inflows between event happening and exhaustion of the investment fund

Figure 1 Structure of the Pool Re funding scheme

By way of example, in 2017 there were three major acts of the terrorism in the UK, the Westminster Bridge attack, the London Bridge attack, and the Manchester bombing at an Ariana Grande concert. The Manchester bombing was the only one of the three where Pool Re was required to pay out, as it involved not only a large amount of commercial damage, but also the

¹⁷ The definition of terrorism can be found in the Terrorism Act 2000.

<https://www.legislation.gov.uk/ukpga/2000/11/section/1>

cancellation of the tour, which falls under the category of Business Interruption caused by damage.

Over the life of Pool Re, they have been involved in thirteen different claims, and have not yet had to turn to the government as a last resort.

Since April 2018, cyber threats have been included in the coverage offered by Pool Re. Previously all risks to property caused by “electronic” means were explicitly excluded from the policies. Since a terrorist act can now be performed remotely via a computer network, and cause property damage, the industry recognises that it is necessary to include this in their coverage. What is excluded from this is the kind of damage which could be perpetrated by low-level hackers – whose intentions are malicious but not political.

Also excluded are acts of war. In 2010 the Stuxnet computer worm was uncovered, which had been developed by the U.S. and Israel with the goal of damaging Iran’s nuclear centrifuges at the Natanz refinery.¹⁸ Damage caused by such a cyber attack would not be covered as it was performed by a state actor.¹⁹

The history of insurance is thus one of tentative steps towards understanding of risk, and the slow development of methods to cover it. Throughout the development of the industry it has been forced to adapt to new social, economic and technical realities. What has been constant is the recognition that insurance is a vital lubricant if the engines of new enterprises are to keep turning.

Globalisation has made the world a smaller place. A side-effect of this is that catastrophic events such as natural disasters and pandemics affect more people and property than ever before. In the next chapter we will look at how the insurance and reinsurance industries have responded to these catastrophes.

¹⁸ <https://arstechnica.com/tech-policy/2012/06/confirmed-us-israel-created-stuxnet-lost-control-of-it/>

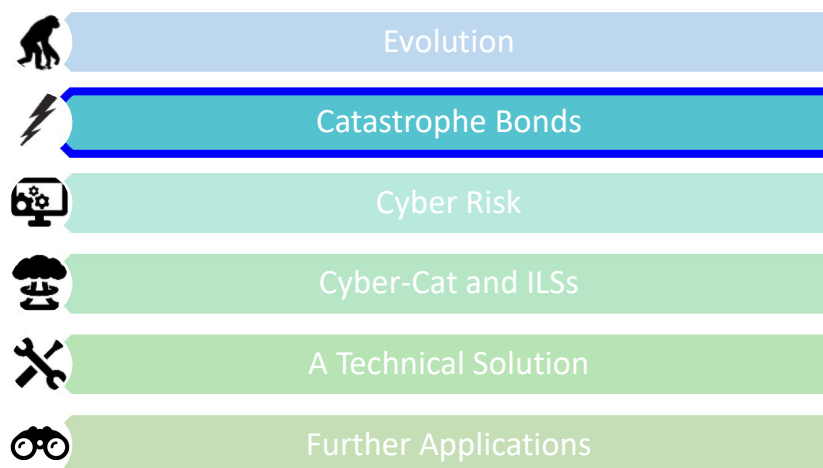
¹⁹ In 2010, the worm escaped into the wider world. Soon more than half the computers in Iran were infected with it. See <https://www.symantec.com/security-center/writeup/2010-071400-3123-99>

2. Insurance-Linked Securities & Catastrophe Bonds

When discussing risk in the commercial sense, one sometimes feels that there are two very different ways of looking at it. On the one hand, there is the insurer, who sees risk as something to be efficiently covered. The losses which might be incurred are skilfully worked out, the premiums, minima and maxima are then backed out from those losses.

The trader, on the other hand, while seeing the risk as something to be priced, is allowed to take a view on it. He can be risk-seeking or risk-averse, he can speculate, he can diversify. What determines his appetite for the species of risk in question can be a variety of different things.

In recent years these two worlds have increasingly intersected, as the insurance industry has looked to offload some of their risk to traders looking for new markets. This has led to the creation of the Catastrophe Bond.



A. History & Motivation

As mentioned in the last chapter, 1992's Hurricane Andrew was a huge blow to the insurance industry, with widespread bankruptcies in the sector. Many firms who had previously been offering reinsurance to cover catastrophes withdrew from the market, leaving many Florida policyholders with either with prohibitively expensive coverage, or none at all.

The government moved to plug some of this gap, but there was still a large hole to be filled. The insurance industry turned to the financial markets to offload some of its risk. The catastrophe bond, or “Cat Bond” was born.²⁰

B. Structure

Catastrophe bonds are now a standard way for a reinsurer to transfer risks to the capital markets. If a catastrophe does occur, the reinsurer can be sure that their payout to the insurance companies is guaranteed.

Technically, Cat Bonds are a subset of Insurance-Linked Securities, or ILSs. However, since most ILSs are linked to the risk of natural catastrophe, the terms “ILS” and “Cat Bond” tend to be used interchangeably.

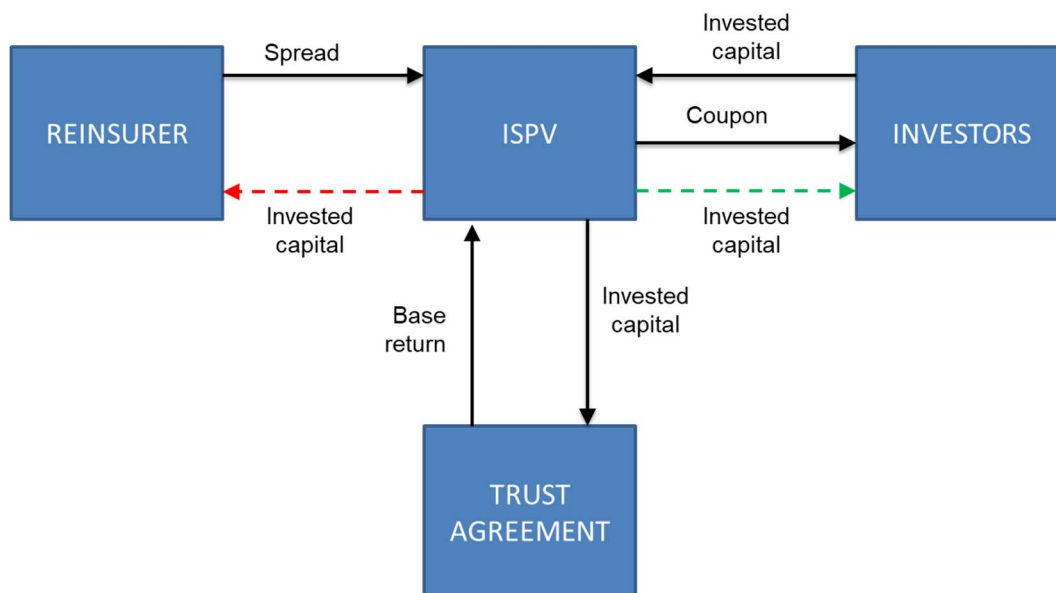


Figure 2 The structure of an ILS

As shown in Figure 2, an ILS is structured as follows:

- The reinsurer sets up an Insurance Special Purpose Vehicle (ISPV) and uses it to issue the Cat Bond.
- Investors buy this bond.
- The ISPV invests the capital in a low-risk market so that it receives some kind of base return.
- The investors receive a coupon from their bond, above the base rate.

²⁰ See <http://en.entropics.se/cat-bonds/the-history-of-cat-bonds/> for more.

- At the end of the life of the bond (generally about three years), the investors get their money back, if no catastrophe has occurred (see the green dotted line).
- However, should the catastrophe occur, the part or all of the capital reverts to the reinsurer who pay out the insurance companies (see the red dotted line.)²¹

The ability to issue Cat Bonds allows a reinsurer to cover catastrophe risk comfortably, while limiting their own exposure. Crucially, reinsurers are permitted to use their ILS issuance to satisfy capital-adequacy regulations.

With reinsurance available, insurance companies also are able to offer coverage to their own customers – ordinary people trying to insure their homes against flooding and hurricane damage. Claims from ILSs can be paid in full in as few as 10 days, and funds can be used as needed for any purpose.

In addition, the arrangement provides a net good to the government. When catastrophe strikes, the government often functions as an insurer of last resort. Cat Bonds ensure that there is a financial buffer in place to absorb losses, and minimise the cost borne by the taxpayer.

In a market dominated by low interest rates ILSs are an attractive proposal for investors. The key appeal of an ILS for an investor is that, in the case of natural catastrophe bonds, the risk is entirely uncorrelated with other market risks. The investor thus has an opportunity to diversify their holdings away from standard capital markets – an appealing prospect over the last ten years. Depending on an investor’s appetite for risk, the returns can be very good – into double figures in some cases.


Figure 3 below, provided by the ILS data authority Artemis, shows that 6.3% of Cat Bonds currently outstanding in provided returns of over 10%.²²

²¹ See

https://www.rims.org/Session%20Handouts/RIMS%2016/RIF003/RIF003_RIF003%20CAT%20Bonds%20101Mon.pdf for more.

²² http://www.artemis.bm/deal_directory/cat_bonds_ils_by_coupon_pricing.html

Catastrophe bonds & ILS risk capital outstanding by coupon pricing




Source: www.Artemis.bm Deal Directory

Figure 3 Outstanding ILS risk by coupon pricing. Source: Artemis. Accessed 27th Oct 2018.

C. Triggering And Parametrisation

As classified by Artemis, most of ILS issuance is dominated by “multi-peril”, i.e. bonds triggered by one of more than one events. An example of multi-peril might be an ILS which could be triggered by either earthquakes or storms. See Figure 4 for an illustration of outstanding risk capital, broken down by peril.

Catastrophe bonds & ILS risk capital outstanding by risk or peril



Source: www.Artemis.bm Deal Directory

Figure 4 Outstanding ILS risk by peril. Source: Artemis. Accessed 27th Oct 2018.

Catastrophe bonds are carefully written to be triggered by defined parameters. Only when very specific conditions are met can the capital revert to the issuer, and be used to cover catastrophic losses.

The ILS can be set up to be on a per-occurrence, aggregate or “only multi-loss” basis. The latter is a structure where the investors do not lose their investment the first time a loss event occurs, but only when more than one event has occurred.

In addition, the nature of the trigger varies.²³ It can be:

- a sliding scale indexed to the dollar losses experienced by issuing reinsurer (“indemnity”),
- a trigger which fires when losses hit a certain industry-wide loss level (“industry loss trigger”), or
- a trigger linked to an index of measured weather or disaster conditions, where measurements above a certain level cause the trigger to be fired (“parametric index trigger”).

23

https://www.rims.org/Session%20Handouts/RIMS%2016/RIF003/RIF003_RIF003%20CAT%20Bonds%20101Mon.pdf

The advantage of the last one, the parametric setup, is that the Cat Bond is linked to an inarguable, centrally agreed, trigger parameter. This provides certainty for issuers and customers. If a Cat Bond is based, for example, on the magnitude of a storm as measured by a central authority, the victims will experience none of the delay introduced by any independent assessment of damage. To this point, the CCRIF (formerly the Caribbean Catastrophe Risk Insurance Facility) has said that their parametric Cat Bonds allowed for payouts to be made in only 14 days.²⁴

Parametric insurance is also customisable. The trigger can be structured optimally to suit issuers' and clients' requirements, allowing the nature of the trigger to be reverse-engineered out of a customer's own loss models. The bond can also be chosen to cover a whole property portfolio, or any subsections thereof.

A recent example of a non-weather-related catastrophe bond is the World Bank's Pandemic Emergency Financing (PEF) facility, which is used to transfer pandemic risk to the capital markets. The facility is supported by an ILS, defined by a number of complex triggers, which provides \$320 million of cover to the PEF. The clear parametric nature of the Cat Bond ensures that areas which have been struck with pandemic disease can receive funds quickly.²⁵

The PEF Cat Bond was developed by Munich Re, who created a trigger carefully designed to reflect the level of contagion. The trigger is composed of a combination of statistics, which include: the death toll, the speed of spread of the disease, and whether the epidemic has managed to jump international borders.²⁶

As we will see, this combination of transparently published data and clear parametrisation make the ILS market a clear candidate for the adoption of technologies which use programmatic logic, such as Smart Ledgers.

²⁴ <http://www.artemis.bm/blog/2018/07/04/ccrif-to-stay-true-to-parametric-roots-of-providing-quick-payouts/>

²⁵ <http://www.artemis.bm/blog/2017/08/03/pef-cat-bonds-to-speed-up-emergency-response-to-pandemics-munich-re/>

²⁶ The PEF bond details can be found at http://www.artemis.bm/deal_directory/ibrd-car-111-112/

D. The ILS Market

Issuance

According to Artemis,²⁷ at the time of writing, the size of the ILS market was \$35.3 billion. However, this number only includes the tradable market. Schroders estimates that the tradable portion of the market only constitutes about a third of the total, the rest being OTC contracts.²⁸ If this is correct, the total size of the ILS market is more than \$100 billion.

The market has traditionally been in a state of boom. From the end of the 1990s to 2005 the market for Cat Bonds grew by an average of 25% annually.²⁹ In the first half of 2017, new records were set with \$8.4 billion of ILS issuance. Despite the major hurricanes of 2017 – Harvey, Irma & Maria – which made it one of the worst years on record for natural disasters – appetite for Cat Bonds remains strong.

What are the reasons for this continued appetite?

- The probability of loss is comparatively low. The National Association of Insurance Commissioners (NAIC) has pointed out that since Cat Bonds started, only 3.3% of the bonds issued have resulted in loss of principal to investors, making them a very good bet.³⁰
- Before the financial crisis, investors used asset-backed securities (ABS's) and mortgage-backed securities (MBS's) as diversification tools. These have since become subject to more regulation and bureaucracy. The availability of the ILS allows for some diversification.
- Another advantage of the ILS over MBSs and ABSs is that the interests of issuers and investors are more aligned. When banks packaged up bad mortgages as MBSs and sell them, the investors were at a severe information disadvantage. ILSs, on the other hand, come with safety features which significantly reduce this risk, in the form of “indemnity triggers”, which ensure that investors' money is only eaten into after

²⁷ http://www.artemis.bm/artemis_ils_market_reports/downloads/q2-2018-cat-bond-ils-market-report.pdf

²⁸

https://www.schroders.com/el/sysglobalassets/schroders/sites/ukpensions/pdfs/2016_march_ils_fi_ve_faq.pdf

²⁹ <http://www.artemis.bm/blog/2018/06/18/ils-at-its-strongest-for-years-after-lessons-of-2017-twelve-capital/>

³⁰ http://www.naic.org/cipr_topics/topic_insurance_linked_securities.htm

losses have reached a certain level. In the language of the MBS, the ILS is effectively the top “tranche” of the risk. There is therefore an alignment of interest between issuer and investor.

- The use of the SPV to hold the collateral means that there is much less exposure to the credit risk of the issuer.
- There are fears that climate change could have a negative effect on the market. However, the short-term nature of the agreements seem to mean that longer-term risks do not affect today’s rates. Publicly issued Cat Bonds have a typical term of three years, whereas private deals are typically a year.

Until recently, natural catastrophe bonds (“Nat Cat Bonds”) have dominated the ILS world, especially those triggered by weather-based disasters in the Americas. There is no reason why this should not change. Recently issuers have used the ILS market to offload life, accident, health and other risks. In July 2018 National Mortgage Insurance Corporation issued a \$264.55 million ILS to cover its mortgage insurance risk.³¹

The Secondary Market

Cat Bonds can be traded on the secondary market. Intermediaries exist who bring buyers and sellers together on a matched trade basis, and provide bid-offer spreads. The prices of Cat Bonds on secondary markets are affected by current meteorological or seismological data, seasonal variation, the reinsurance market, and the behaviour of the rest of the capital markets. If other markets are experiencing tight liquidity, Nat Cat Bonds may experience a boost as people turn to an uncorrelated market.

The models used to price Cat Bonds in the secondary market are similar to the stochastic models used in other markets. A popular mechanism is to use a “Poisson process”, which is a mathematical function used to model events which arrive at random times. It is commonly used in finance to model bond default times. In addition, since the Cat Bond is effectively a kind of Floating Rate Note, interest-rate models are incorporated into the pricing calculations. Further terms are added to the models to represent the dollar amount lost given a catastrophic event, calibrated using historic data.³²

³¹ http://www.artemis.bm/deal_directory/oaktown-re-ii-ltd/

³² See <http://www.hdbresearch.com/index.php/hdbr/article/download/64/79/> for an overview of past work.

The popularity of the secondary market is such that Extraordinary Re, in partnership with Nasdaq, is launching a new platform to allow investors to trade exposure to insurance risk.³³ This will make it much easier for the average investor to take part in what has historically been a relatively opaque market.

As the secondary market volumes expands, the sophistication of the financial mathematics used in pricing will only increase. An indicator of the increase in interest in this market can be seen in the EurekaHedge ILS Advisers Index, which shows increasing levels of return over the last twelve years - with a significant dip in 2017, thanks to the unusual amount of hurricane activity in the Gulf of Mexico.³⁴

E. London's Changing Role

In December 2017, the UK passed legislation which allowed London reinsurers to undertake ILS business in the UK. Previously that business had to be undertaken in Bermuda. Within two weeks of the legislation, the Neon Group took part in the first ILS deal.³⁵

Now that ILSs can be written in London, the market is set to grow faster than ever. London's differentiating factors include easy access to the capital markets, underwriting expertise, experience in issuing and trading complex financial securities, and swift approval of SPVs. Another advantage is that there are many existing restrictions on governments which prevent them from investing in off-shore schemes. The door is now open to investment from the U.S. and the European Union.

F. Terror-Linked ILS

As discussed in the previous chapter, the Pool Re group was formed in order to provide a government-sponsored backstop for insurance against property damage caused by terrorism.

³³ <https://www.reuters.com/article/us-usa-exchange-insurance/startup-extraordinary-re-to-launch-exchange-for-insurance-risk-idUSKCN1GD4XE>

³⁴ http://www.eurekahedge.com/Indices/IndexView/Special/635/Eurekahedge_ILS_Advisers_Index

³⁵ <http://www.artemis.bm/blog/2018/01/03/neons-uk-ils-vehicle-ncm-re-enters-into-72m-syndicate-quota-share/>

Since 2016, Pool Re have been doing a feasibility study into using ILSs to support terrorism insurance.³⁶ With the exception of a small bond used to insure against cancellation risk at the 2006 World Cup in Germany,³⁷ there terrorism. Prospective customers are understandably skittish. In order to feel comfortable with the market, they need not only understand the deterministic element of modelling loss amounts, they must also be comfortable with the probabilistic aspect of the problem – predicting the likelihood of an attack in the absence of data.

One bright spot is that despite the lack of an existing ILS market, investors can look to the existing non-ILS retrocession market for assurances. Pool Re pointed out to us that if a very conservative reinsurer such as Munich Re is involved in buying and selling retrocession on terrorism risk, the investor can take the view that the market is perhaps not too unsafe.

Pool Re report that their research is going well so far. They have investor interest and are considering an issue of about £100 million, to support terrorism insurance. The standard retrocession renewal period for the reinsurance market comes around in spring 2019, and Pool Re hope to issue the ILS at that time.

So far the proportion of non-property-catastrophe ILSs sold is miniscule, as was shown in Figure 3. Many analysts believe that there is a great deal of scope outside property for the ILS market to expand into new areas.³⁸

In the next chapter we will look at the most exciting of these new areas: cyber risk.

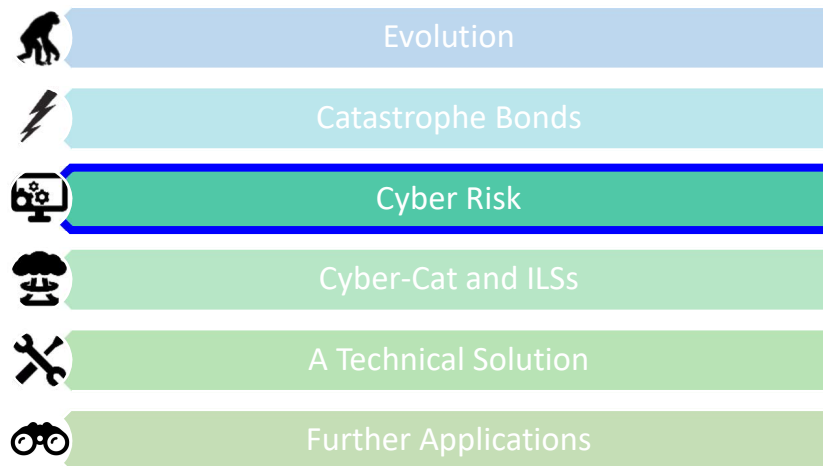
³⁶ <http://www.artemis.bm/blog/2018/04/25/pool-re-explores-terror-ils-options-with-help-of-gc-securities/>

³⁷ <http://www.rms.com/blog/2014/07/16/rms-and-the-fifa-world-cup-insuring-against-terrorism/>

³⁸ <http://www.artemis.bm/blog/2017/07/21/ils-cat-bond-growth-strong-but-one-dimensional-johansmeyer-pcs/>

3. Cyber Risk

Before considering how ILSs could be applied to cyber-catastrophe, it is worth considering how the insurance industry treats cyber risk in general at present.



A. Definition

The term “cyber”, as used by the insurance industry, is extremely broad. It covers everything from actual theft, business interruption, failure of hardware or software, to regulatory fines imposed in the wake of data loss. It seems to the outsider that any risk which remotely involves the use of computers or networks is considered to fall under the cyber umbrella.

Some of the risks which are deemed to be cyber risks in the literature³⁹ are:
Theft, including:

- Mass theft of credentials
- Data espionage
- Financial fraud
- Cash theft

Disruption, including:

- Power grid disruption
- MS Windows exploit
- Tran systems disruption
- Comms silenced

³⁹ https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/crs-sybil-logic-bomb-cyber-catastrophe-stress-test.pdf

- GPS failure
- Tactical data espionage
- Degrading of internet and denial of service

...and Damage, including:

- Long term data corruption
- Leaks, abuse of data, defamation
- Data centres, internal IT and cloud servers damaged
- Targeted physical damage
- Algorithmic systems failures

This definitional problem has been widely discussed. As mentioned by Biener et al (2015),⁴⁰ many attempts have been made to pin down exactly which risks qualify. Some authorities only admit risk of malicious damage as a cyber risk, others take an extremely broad view, classing anything and everything related to information security as a “cyber”. Biener et al themselves focus on “operational risks”, which have historically been the risks which insurers themselves are interested in.

Definition is important. Insurers need to be able to write sensible, well-defined policies, so that they, their customers, and their reinsurers are clear on exactly what is covered and what is not. At the moment, not only is the definition of cyber risk rather broad, the coverage available under the name “cyber” is often patchy. That is, not only are the risks themselves unclear, but many of them are implicitly or explicitly excluded from many policies without customers being aware of it. On the other side, from the insurers’ point of view, the policies may involve a great deal of unforeseen silent cover and aggregations. If these could be identified, they could be explicitly excluded from the policy and sold as standalone products or endorsements, to the benefit of both sides.

40

www.alexandria.unisg.ch/238242/1/15_03_Biener%20et%20al_Insurability%20of%20Cyber%20Risk.pdf

B. Current Insurance Offerings

In 2016, the Judge Business School found that across the international market, there were a wide variety of cyber loss types included in cyber insurance products, as shown in the table below.⁴¹

Cyber Coverage	% of Products Offering this Cover
Breach of privacy event	92%
Data and software loss	81%
Incident response costs	81%
Cyber extortion	73%
Business interruption	69%
Multi-media liabilities (defamation and disparagement)	65%
Regulatory and defense coverage	62%
Reputational damage	46%
Network service failure liabilities	42%
Contingent Business Interruption	33%
Liability – Technology Errors & Omissions	27%
Liability – Professional Services Errors & Omissions	23%
Financial theft & fraud	23%
Intellectual property (IP) theft	23%
Physical asset damage	19%
Death and bodily injury	15%
Cyber terrorism	12%
Liability – Directors & Officers	13%
Liability – Product and Operations	8%
Environmental damage	4%

As can be seen, most of the insurance taken out under the cyber umbrella is in fact to cover breach of privacy costs: the administrative expenses associated with notifying all users of the data loss, and the costs of fighting possible resultant legal action. Still, Business Interruption (BI) coverage is greatly in demand. 69% of cyber insurance products include BI coverage of some kind.

⁴¹ Taken from

https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/crs-rms-managing-cyber-insurance-accumulation-risk.pdf

The market for cyber-related physical damage to operational tech is much less in demand. Interference in, or failure of, the control systems for physical processes in factories or power stations could be exceedingly financially damaging and cause loss of life, and yet there is little demand for coverage.

A welcome development in the thinking on cyber exposure has been the publication of The Cyber Insurance Exposure Data Schema.⁴² This is a result of collaboration between Judge Business School at Cambridge and multiple insurance providers. The aim is to assist insurers and insureds to clarify what their policies cover. The schema document points out that there are four categories of cyber exposure cover:

- Affirmative stand-alone cyber cover in the form of specific policies for data breach, liabilities, property damage, and other losses resulting from information technology failures, either accidental or malicious.
- Affirmative cyber endorsements which extend the coverage of a traditional insurance product, such as commercial general liability, to cover cyber-induced losses, typically to cover a privacy breach.
- Silent cyber exposure owing to gaps in cyber exclusions. A policy may have exclusion clauses for malicious cyber-attacks, apart from certain nominated perils, for example, Fire, Lightning, Explosion, and Aircraft Impact (FLEXA). The policy therefore has exposure to a cyber-attack if one were to trigger one of the nominated perils to cause a loss.
- Silent cyber exposure from policies without any cyber exclusions. Many insurance lines of business incorporate 'All Risks' policies without explicit exclusions or endorsements for losses that might occur via cyber-attacks.

The Schema defines the Business Interruption category of loss as *"Lost profits or extra expenses incurred due to the unavailability of IT systems or data as a result of cyber-attacks or other non-malicious IT failures."* A key aspect of this definition is that it can be malicious or not. If the business is interrupted the insurer covers the cost.

One thing which distinguishes BI from other types of cyber risk is that the event is immediately detectable. In the case of data theft or a security compromise, it may take months or even years to fully calculate the extent of the losses, in

⁴² https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/crs-cyber-data-schema-v1.0.pdf

order to place a claim. A cloud or network outage causes immediate, calculable losses. A policyholder may need an immediate infusion of cash in order to stay afloat.

In their research on accumulation risk,⁴³ the Judge Business School group identified five cyber loss processes:

- Cyber Data Exfiltration
- Denial-of-Service
- Cloud Service Provider Failure
- Financial Transaction Cyber Compromise
- Cyber Extortion

Of the five cyber loss processes above, the two which would lead to immediate BI are Denial-Of-Service and Cloud Service Provider Failure.

C. Attacks Causing Business Interruption

DoS And DDoS

A Denial-of-Service (DoS) attack is where a server is flooded with a massive amount of junk information, rendering it inoperable. If the attack comes from more than one source – in other words the malicious party has set up multiple sources of junk data – then the attack is classified as a Distributed Denial-of-Service (DDoS) attack.

If an attack sends the junk data at speeds of 1-10 Gbps (gigabits per second), this is termed a “significant intensity” attack. At more than 100 Gbps, the attack is deemed “very high intensity”. In 2014, a total of 1.6 million DoS attacks were recorded, of which 500 were classified as very high intensity.

How much data is required to shut a server down depends on the amount of traffic the server is designed to support. For example, a website intended to support a million users per month can be overwhelmed by an attack of 10 Gbps or above. Attacks of greater than 600 Gbps have been recorded.

⁴³ “Managing Cyber Insurance Accumulation Risk”,
https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/crs-rms-managing-cyber-insurance-accumulation-risk.pdf

Most attacks are short in duration. 50% of them last for less than two hours, and 70% are over in under six hours. 16% of them last longer than twelve hours, a number which is significant for insurers, since this is the most common deductible period for Business Interruption insurance policies.

In 2014, there were 500 DDoS attacks which were of either high or very high intensity and lasted for longer than twelve hours.

DoS attacks fall into three categories:⁴⁴

- Volumetric: Attacks which use massive amount of traffic to saturate the bandwidth of the target.
- Application-based: Attacks which establish a connection with the target and then exhaust the server resources by monopolizing processes and transactions.
- Protocol-based: Attacks which consume all the processing capacity of the target, or intermediate critical resources like a firewall.

Application and protocol-based attacks are the most sophisticated and can be the most challenging to identify.

The motivations for the attacks are usually simply destructive. They are performed as a form of protest against large companies, using sabotage and vandalism to affect the businesses. Some attacks are used to extort money, and some are used as a distraction to draw attention away from other, more damaging activities, such as stealing data or IP.

To defend against these attacks can be difficult. The most common defence strategy is to analyse all traffic coming into the system, and filter out the junk while letting the real messages through. This is often done at the ISP level, although large corporations will have their own firewalls and filters.

In the event of an attack breaching these defences - if say, the junk is efficiently disguised as valid messaging, it can take time for new strategies to be put into place. Backup or disaster recovery systems will need to be brought up and synced with the main systems. The recovery and repair time may therefore lead to significant Business Interruption.

⁴⁴ <https://blog.thousandeyes.com/three-types-ddos-attacks/>

For example, many attacks have been performed by a tool called the Great Cannon of China, which assists the Chinese government in their internet censorship goals.⁴⁵ This tool is capable of intercepting innocent web traffic travelling to or from specified servers, altering the data and redirecting it at target servers. In March 2015, some of the traffic sent to the search engine Baidu was altered to contain a malicious request script. This script was then pointed at GitHub, the code-sharing website. The volume of requests proved too much for GitHub to handle, and the domain was intermittently affected for a number of days.

In cases like this, when state actors are involved, insurers often refuse to pay out. Most, if not all policies have exclusions in the case of war, which in practice means any activity performed by a sovereign state.

In 2012, the “hactivist” group known as “Anonymous” announced their intention to launch a DDoS against the Internet’s thirteen root domain name system (DNS) servers. They called this plan “Operation Global Blackout”. It came to nothing, but at the time, most experts agreed that the plan as described by the group was feasible, if difficult. Since then, measures have been put into place to make such an attack harder, but it is still possible in principle given sufficient ingenuity.

Indeed, in October 2016 a DDoS attack was aimed at the servers of the DNS host Dyn, causing problems at multiple websites such as Twitter, Reddit, Spotify, the New York Times.⁴⁶

Cloud Service Outage

The use of cloud services has exploded in recent years. Many businesses find that significant savings are to be made by running their software “in the cloud” as opposed to on-site. A Logic Monitor survey conducted in 2017 concluded that by 2020, 83% of Enterprise Workloads will be in the cloud.⁴⁷

⁴⁵ <https://www.theguardian.com/technology/2015/apr/13/great-cannon-china-internet-users-weapon-cyberwar>

⁴⁶ <https://www.technologyreview.com/s/602709/massive-internet-outage-could-be-a-sign-of-things-to-come/>

⁴⁷ <https://www.logicmonitor.com/resource/the-future-of-the-cloud-a-cloud-influencers-survey/>

The cloud provision market is dominated by a few large providers. The market leader, Amazon Web Services (AWS), have a 33% share of the market, trailed by Microsoft at 13%, and IBM, Google and Alibaba all at well under 10%.⁴⁸

If a large number of disparate businesses are dependent on a small number of providers, it follows that the risk of a large-scale outage increases, affecting multiple businesses. A recent study by Lloyd's of London⁴⁹ classified the possible causes of cloud outage into the following vectors:

- Environmental – this includes natural causes like lightning strikes and earthquakes, as well as terrorist attacks – anything that could shut a data centre down through physical means.
- Adversarial – this includes a DDoS on the cloud service provider, or a deletion of the cloud's virtual machines by a malicious insider.
- Accidental – this is mostly concerned with internal IT mistakes, such as backup failures or human error.
- Structural – failures in power systems, networking devices or disks.

Of the above, the last two have led to cloud downtime. The longest outage recorded was a three-and-a-half day outage of AWS in 2011, when a transformer failed. This was severe, but fairly localised, and did not affect all the cloud service's users. However, in both February and October 2013 Microsoft Azure suffered worldwide outages.

The Lloyd's report indicates that if AWS were to experience an outage of between twelve hours and a full day, the total loss incurred by users would be \$5.89 billion. Of this loss, about \$4.83 billion is technically insurable, but only \$1.08 billion is currently insured. At the extreme end, for outages of between 5.5 and 11 days, the gross loss would be \$23 billion, of which \$19.91 billion is insurable, but only \$4.32 billion is insured. These figures may in fact be low, since they assume a "waiting period" of 8 hours before losses begin accruing.

An earlier Lloyd's study⁵⁰ points out vulnerabilities inherent in so-called "hypervisor" software. Hypervisors allow a cloud's clients to share software operated by the host, while still remaining separate. As far as the client is

⁴⁸ <https://www.cnbc.com/2018/04/27/microsoft-gains-cloud-market-share-in-q1-but-aws-still-dominates.html>

⁴⁹ <https://www.lloyds.com/~media/files/news-and-insight/risk-insight/2018/cloud-down/aircyberlloydspublic2018final.pdf>

⁵⁰ <https://www.lloyds.com/~media/files/news-and-insight/risk-insight/2017/cyence/emerging-risk-report-2017---counting-the-cost.pdf>

concerned, they are using a single machine, but in fact the resources are being shared among the cloud's customers.

Two popular hypervisors are:

- Xen, used by Amazon, Rackspace and IBM. This has suffered from code bloat (now over 1.1 million lines) and in 2016, a review discovered 28 public vulnerabilities.
- ESXi, which was created for the popular system VMWare, used by many Enterprise clients. Multiple security flaws have been found in the last few years.

When vulnerabilities are discovered, remedial action often necessitates a reboot of the hypervisors, which results in outages. As the 2017 Lloyd's report says:

“Problems with hypervisors can lead to large-scale issues. For example, in 2015, Amazon AWS had to reboot its systems on two occasions due to Xen hypervisor patches that required a full restart of all affected systems. The software flaw, CVE-2015-7835, affecting the Xen hypervisor allowed malicious actors to create guest servers which could access the host computer's memory and take control of the entire system. This security breach went undiscovered for more than seven years.”

Hardware

When discussing the possibility of Business Interruption, it is natural to focus on the network, which is hard to control centrally, and software, which is relatively easier to alter maliciously. Hardware, while never infallible, tends to be regarded as something which is at the mercy of interruption only from *external* factors such as power outages or lightning strikes.

It has come to light, however, that hardware may not be as neutral as all that. In 2015, one of the U.S.'s largest motherboard manufacturers, Super Micro Computer Inc., discovered that their Chinese manufacturers had inserted on some boards a microchip, the size of a grain of rice. These chips allowed hackers to get into any network which included the altered machines.⁵¹

⁵¹ <https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies>

If altered boards such as these became sufficiently widely used, it would be possible to cause vast network outages, perhaps far bigger than any DoS attack could accomplish.

D. Current Reinsurance Capacity

2017 was a particularly bad year for the insurance industry. Not only were there three major hurricanes in the Caribbean and the Gulf (Harvey, Irma, Maria), there were also six California wildfires, causing massive property damage. Insured losses in the U.S. for the second half of 2017 reached \$70 billion. Following this, a few major cyber-catastrophe events occurred in the last quarter of the year, including the NotPetya virus. This affected hundreds of businesses, in particular the shipping company Maersk, who reported losses of up to \$300 million, and the pharmaceutical company Merck.⁵² There was also a major data breach at the Equifax credit reference company, causing losses in excess of \$1 billion.⁵³

According to some authorities, these losses represent a major opportunity for insurers.⁵⁴ The new cyber-risk is under-insured, and the ILS market is seen as increasingly competitive. Cyber first-party insurance covering Business Interruption suffers from lack of uptake. Why might this be?

Biener et al. (2015), investigating the insurability of cyber losses, found that there are three main stumbling blocks preventing the market from taking off:

- Highly interrelated losses (thanks to insufficient diversification)
- Lack of data, and
- Severe information asymmetries.

All three of these problems are to some extent lessening. It has been found that as more participants enter the U.S. market, premiums have decreased somewhat.

Lack of data has a serious impact on the price of coverage. Without data, risks are hard to estimate, so insurers tend to set high deductibles and low maximum

⁵² <https://www.verisk.com/siteassets/media/pcs/pcs-global-risk-loss-report-2017.pdf>

⁵³ <https://www.reuters.com/article/us-equifax-cyber/equifax-breach-could-be-most-costly-in-corporate-history-idUSKCN1GE257>

⁵⁴ <https://www.casact.org/community/affiliates/camar/0518/Johansmeyer.pdf>

coverage. The lack-of-data problem is being tackled head-on. Legislation has been passed so that reporting of events is now compulsory.

Other problems cause sluggishness in insurance uptake. A recent EU report⁵⁵ made mention of the fact that not only is data lacking, preventing anyone from making adequate risk assessments, there is also a lack of cybersecurity skills in the cyber insurance market. This results in carriers frequently having limited understanding of the risks, or what data is needed in order to perform the assessment.

As for information asymmetries, these can be mitigated to some extent by an upfront risk assessment of a client's systems, using a reliable standard such as the Security Effectiveness Score.⁵⁶ A typical length of time for cyber policy is just one year, allowing for regular security re-assessments.

However, while upfront risk assessments can help with security questions, the benefits they bring are probably not going to be of any use in a DDoS scenario. Catastrophic network outages can only be addressed by putting in place a reliable financial mechanism to cover clients quickly and transparently.

E. Current Cyber Insurance Technology

In recent years a number of new and existing technology companies have entered the cyber-insurance market.⁵⁷ These new players provide a cyber-risk modelling service to the insurance industry, chiefly underwriters and brokers. Some examples are:

- ThreatInformer. Based only on freely available external data, they can produce for an underwriter “a unique cyber insurance profile for each insured.”
- At-Bay, which sells cyber insurance and consultancy to clients, as well as providing analytics for insurers.
- Cyence, which has a special focus on catastrophe risk. They collect a great deal of data from the internet and other external sources, and apply machine learning techniques to filter the data and detect relevant patterns.

⁵⁵ <https://www.enisa.europa.eu/publications/commonality-of-risk-assessment-language-in-cyber-insurance>

⁵⁶ <https://blog.focal-point.com/measuring-security-and-the-financial-impact-of-data-breaches>

⁵⁷ <https://techcrunch.com/2016/05/23/can-startups-disrupt-the-20-billion-cyber-insurance-market/>

These types of software solutions are welcome, and contribute to a wider understanding of cyber risk in general. However, these solutions are all focused on security.

A Business Interruption cyber-risk catastrophe, when it comes, will be like a power black-out – instantaneous and wide-ranging. The probabilistic security models arrived at beforehand are not of any use in the case of a single catastrophic event.

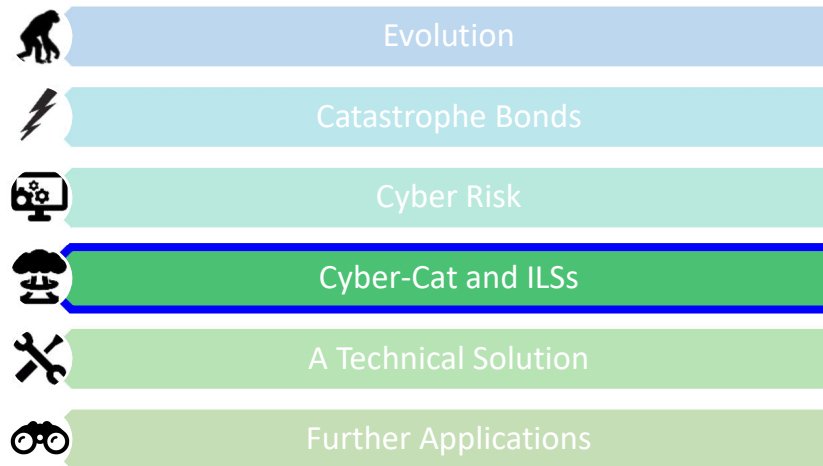
What is lacking is not modelling and data gathering, but a joined-up mechanism to monitor the larger system in real time and connect gathered system data to reinsurers. Such a mechanism should be able to firstly detect the exact extent and nature of the outage, and trigger financial payments – within days or hours – to those who need it.

When a catastrophe occurs the insurance industry needs a cash injection quickly - perhaps within such a short period that the insurer cannot investigate the cause of the event in time. In the case of a cyber-catastrophe, the insurer may have to resign themselves to treating the “climate” of a computer network like the Earth’s climate - in other words, accepting that events may not have an easily findable root cause. The determination of the *cause* of the outage should not delay payments, just as the origins of natural catastrophes are not pored over post hoc. The detection and measurement of the outage alone should be sufficient to trigger payments.

In the next chapter we discuss how this type of “parametric” approach is helpful when using ILSs to cover cyber-catastrophe risk.

4. Cyber-Catastrophe And ILS

In this chapter we bring together the concepts of cyber-catastrophe and Insurance-Linked Securities. How might financial markets help to manage the risk of massive cyber failure?



A. A Gap In The Market

As we have seen, the world is facing the possibility of a large systemic network outage which may strike at any time. Either through malice, carelessness or sheer accident, it is possible, if unlikely, that a wide-scale network failure could result in significant Business Interruption losses. Many of these losses are drastically understated.

For some time, the Financial Times has been highlighting this large gap in the market.⁵⁸ Analysis suggested that some institutions might require cyber insurance support of as much as \$1 billion, and appropriate policies were simply not available. For example, when the U.S. retailer Target suffered a major data breach, it affected 110 million of its customers. The overall cost of the breach was estimated at \$248 million, of which only 36% was covered by Target's insurance.⁵⁹

These gaps exist in all branches of cyber insurance, from Business Interruption to data breach liability cover. Given these gaps, it is clear that in the case of a

⁵⁸ <https://www.ft.com/content/61880f7a-b3a7-11e4-a6c1-00144feab7de>

⁵⁹ <https://www.insureon.com/blog/post/2015/03/24/how-much-does-your-cyber-liability-insurance-cover.aspx>

large cyber-catastrophe the government will have to act as the insurer of last resort by default.

A number of commentators believe that given the dangerous amount of cyber risk which firms are knowingly or unknowingly exposed to, it is a mistake to wait any longer for a Cyber-Cat Bond to be issued.⁶⁰

Many feel that far from being an insoluble problem, this market provides a huge reserve of untapped opportunities for the insurance industry. There is a high demand for new avenues of investment. Since the introduction of the natural catastrophe ILS in the 90's, the uptake has been slow and steady. We have now reached a point that the market is mature, with diminishing returns available. Investors are looking for more diversification, particularly in a low-interest-rate environment.

B. More Data More Models

What is preventing potential buyers, with appetite for risk, from meeting sellers, with risk to offload? Clearly, buyers feel that their risk-based concerns are not being addressed.

A major obstacle standing in the way of such a cyber-catastrophe bond issuance is the perceived lack of data. In the EU, this problem is beginning to be addressed. The EU has made it mandatory for institutions to report on cyber-attacks. The Association of British Insurers (ABI) has appealed for a national database of information about cyber incidents, to which everyone affected by an incident would have to submit incident details.

One of the reasons that cyber-catastrophe is intimidating from a risk perspective is that the "big one", i.e., the cyber equivalent of Hurricane Andrew, has not yet happened. Nobody quite knows how bad the fallout from a mass cyber event could be.

In addition, cyber-catastrophe risk is seen as very hard to model. The perception is that the technology, data and modelling do not exist or are insufficiently developed. This is not the case. The required risk and data methodologies are there, but not yet widely understood.

⁶⁰ <http://www.artemis.bm/blog/2017/02/08/stop-waiting-for-a-cyber-cat-bond-johansmeyer-welsh/>

As mentioned above, the publication of the Cyber Insurance Data Schema by the Judge Business School has provided a robust cyber underwriting taxonomy. It is now down to the industry to attempt to deliver robust analytics.

The modelling of both the risk of a cyber-catastrophe event, and of the losses due to the event, need to be developed. It is thought that network outages could follow existing models of disease contagion.⁶¹

Using pandemic-based models, the outages could be more tractably simulated in order to work out the risk and extent of losses. Like the development of all models, this will involve a trade-off between model complexity and computational efficiency.

C. Structuring The Risk

When risk is bought and sold in the capital markets, data is often in short supply. In such situations, the seller of the risk can frequently work wonders by means of appropriate structuring of the security. If a risk can be sliced up and packaged in an appealing way, with well-defined parameters, it not only reduces the investor exposure, it also makes the security much easier to model.

In 2015 an Artemis article discussed the possibility of doing just this. Could the reinsurance industry offload their cyber-security risk on to the capital markets with the help of effective structuring?⁶²

The obvious solution would be to apply existing catastrophe-bond structures to the cyber problem. However, the current structure of the catastrophe-bond may be putting investors off. The number of unknowns perceived to be lurking in the investment do not inspire confidence, especially since Cat Bonds demand money up front.

The 2015 article suggests that one way around this problem could be the use of contingent capital. In this arrangement the investors would effectively promise to pay out the full amount when the structure is triggered.⁶³

⁶¹ <https://www.soa.org/Files/Research/Projects/cybersecurity-insurance-report.pdf>

⁶² <http://www.artemis.bm/blog/2015/03/23/could-the-capital-markets-solve-the-1b-cyber-insurance-policy-gap/>

⁶³ <http://lexicon.ft.com/term?term=contingent-capital>

The problem here is the inevitable introduction of credit risk. In addition, the distributed nature of the obligation would almost certainly result in payments being made piecemeal and tardily at that. For the insurance and reinsurance industries looking for certainty, this option may be too unreliable.

D. Triggers

The ILS community is wary of catastrophe risk partly as a result of the lack of defined coverage limits and triggers. If these are brought to bear on the problem, capacity from the ILS community could be forthcoming.

Johansmeyer and Welsh (2017) maintain that it is perfectly possible to render the modelling of risk tractable.⁶⁴ They say that it is currently possible, with the available analytics, for parametric triggers to be developed and marketed.

Parametric insurance has perhaps more in common with the world of capital markets than the world of traditional insurance.⁶⁵ Someone will buy protection based on a defined event, rather than the loss they incur. A parametric trigger is a mechanism to pay out based on inarguable metrics. The risk associated with these metrics will be related (but not identical) to the risk the buyer wishes to acquire coverage for. Fully deterministic triggering events would minimise the risk of loss aggregation from “silent” coverage.

Triggers are the key to sorting out another drawback in the current cyber insurance world – that of time to settle. Just as in natural catastrophes, quick payment is required to allow frontline insurers to pay their customers. Fully deterministic metrics are ideally suited to this problem, since it is easy to work out whether the event is triggered or not, and make the payment.

Other commentators⁶⁶ have maintained that the use of proportional agreements may be key in the allowing the market to grow. In the ILS market today, most deals are “excess of loss” types, where a certain amount is paid or not depending on the outcome of a claim. Proportional agreements will facilitate the designing of more complex triggers to match the appetite of investors, and will promote sharing of expertise between the capital markets and the insurance industry.

⁶⁴ <http://www.artemis.bm/blog/2017/02/08/stop-waiting-for-a-cyber-cat-bond-johansmeyer-welsh/>

⁶⁵ <http://www.artemis.bm/library/what-is-parametric-insurance.html>

⁶⁶ <http://riskandinsurance.com/cyber-risks-ils/>

Of course, once the structuring Jack is out of the box, the opportunities for bespoke tailoring of the parameters are limitless. If insurers are uncomfortable with “pure” parameterisation, and would prefer that the size of loss incurred by a specific cedent still plays a role, “hybrid” protection can be developed – that is, protection which is not fully parametric, but has a parametric component.

E. New Markets For ILS

As we have seen, the best way for the reinsurance industry to deal with natural catastrophe risk is to transfer it the capital markets by means of an insurance-linked security, or ILS. Many new entrants to the market are waking up to the fact that there is an appetite for cyber-based ILS. Not only that, the ILS structure neatly does away with many of the drawbacks inherent in proportional agreements and contingent claims. Contingent claims suffer from widespread (and possibly unpleasantly correlated) credit risk, since the reinsurer is dependent on timely and reliable payments from multiple participants. An ILS structure solves this problem through the use of a centralised SPV.

Notions of quick and efficient payment triggers are bread and butter to the capital markets, which rely on the swift settlement of futures, options and other derivatives to oil the wheels. ILSs structured in this way are very attractive to investors. It may be that the very presence of trigger-based ILSs will stimulate quantitative research into appropriate risk models for cyber-catastrophe.

Fortunately, there are new entrants to the market who recognise these opportunities. Ridge Global are offering “intelligent” cyber insurance in Canada, with hybrid solutions available.⁶⁷

Other institutions⁶⁸ are offering analytic services explicitly focused on cyber risk modelling, with an emphasis on non-traditional (i.e., non-actuarial) methods.

⁶⁷ <https://www.ridgeglobal.com/david-peterson-and-tom-ridge-partner-to-offer-intelligent-cyber-insurance-in-canada/>

⁶⁸ See for instance CyberCube at <https://www.cybcube.com/>

F. The Future Of ILS Cyber

According to Jean-Louis Monnier, the head of ILS structuring at Swiss Re, the introduction of cyber-catastrophe ILSs is only a matter of time.⁶⁹

What will the impact of this be? How will investors respond to this new asset class?

A BNY Mellon report⁷⁰ points out that one motivation to invest in Nat Cat Bonds is that they are almost totally decorrelated with equity markets. A downturn in the markets does not lead to a natural catastrophe. On the other hand, while a catastrophe may have knock-on effects in the equity markets (as traders pull out of affected businesses) the bounce in equity prices is often not very long-lasting.⁷¹

The BNY report suggests that cyber-catastrophes may be different from natural catastrophes in this respect. The report makes the point that there is (or is perceived to be) a larger correlation between cyber-catastrophes and equity markets. In principle, a cyber-catastrophe has the potential to adversely affect the equity prices of all the businesses involved in the catastrophe. This means that cyber markets could be a less effective form of diversification from standard markets. Would investors therefore be less interested in cyber-catastrophe risk than in natural catastrophe risk?

Market attitudes are hard to predict, but one can make some headway in the analysis by thinking about the nature of the correlation. On one hand, cyber risk is identical with Nat Cat risk in that the correlation with equity markets only happens in one direction. Equity downturns don't cause catastrophes, be they cyber or natural. So, if an investor wants to reduce their exposure to non-catastrophe-related market crashes, any kind of Cat Bond will do. So, on the other hand, the question becomes: is the effect of a natural disaster on equity prices the same as the effect of a cyber-catastrophe on equity prices? In other

⁶⁹ <http://www.artemis.bm/blog/2017/10/04/cyber-cat-bonds-will-be-a-reality-within-two-years-jean-louis-monnier-swiss-re/>

⁷⁰ <https://www.bnymellon.com/emea/en/locale-assets/pdf/our-thinking/insurance-linked-securities-cyber-risk-insurers-and-the-capital-markets.pdf>

⁷¹ https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/170622-slides-mahalingam.pdf

words, would the equity bounce in the wake of natural catastrophe be similar in the wake of cyber-catastrophes?

To answer this, let us look at the effect of the AWS outage in February 2017. While it lasted, this four-hour outage did indeed have an effect of the stock prices of many S&P 500 companies, wiping \$150 million off the index.⁷²

How quickly did it recover? Price history suggests that while online retailers such as Target were hard hit, and did not recover for some time, companies not so directly dependent on the Web, such as Nike, did not experience a great deal of equity price devaluation.

In the case of natural catastrophes, similar logic holds. Equity prices for certain sectors are harder hit than others.⁷³ The sectors most affected are the hospitality industry, and the insurance industry which covers it. Online retailers are not be significantly affected.

So, it is perhaps too general a statement to say that equity prices as a whole are affected by catastrophes. Cyber-Cat Bonds could serve as an effective hedge in cases where natural Cat Bonds are overly correlated with certain sectors, such as bricks-and-mortar retail.

This analysis indicates that there is a definite niche in the market for cyber-catastrophe ILSs. The introduction of such securities can only assist in the effective hedging of risk across currently under-served sectors.

We have established that networked computer systems are vulnerable to catastrophe, and that managing this risk using financial markets would be a good idea. The reader may have spotted a problem with this. Given that financial markets cannot run without networked computers, and that networked computer systems are the very thing which are at risk, how can we ensure that

⁷² <https://www.businessinsider.com/aws-outage-hurt-internet-retailers-except-amazon-2017-3>

⁷³ <https://www.marketwatch.com/story/what-history-says-about-hurricane-irma-and-the-stock-market-2017-09-08>

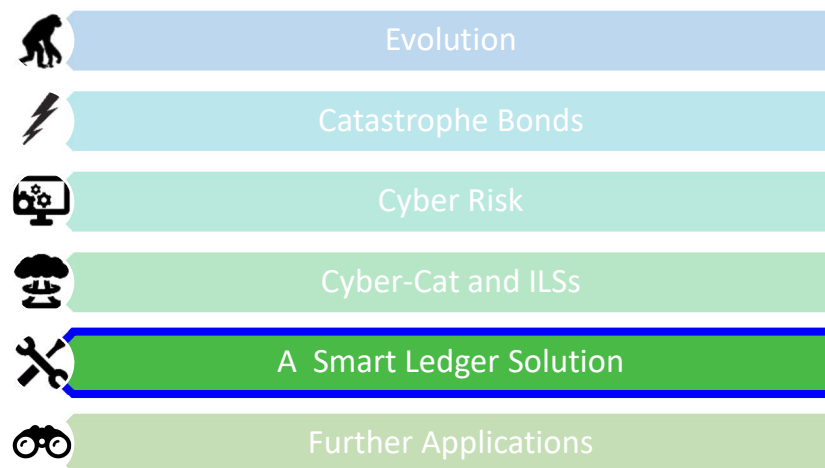
the system is resilient enough to withstand the catastrophe, and ensure that payments are made as required?

In the next chapter we will propose an architecture to help solve this problem.

5. A Smart Ledger Solution

In light of all the foregoing, we can see that there is a clear market for some way to trigger an ILS payment. Our discussion of cyber-insurance generally has identified that covering the risk of Business Interruption caused by a network outage is the sort of cyber-insurance most amenable to this treatment.

In order to allow for the writing of an ILS prospectus with an inarguable set of triggering conditions, what is required is an agreed, unambiguous metric indicating the nature and extent of a network outage. This metric will need to be clearly defined and neutral, so that both the reinsurer issuing the ILS and the investor in the bond will fully understand their exposures.



A. The Network Availability Index

What might such a metric consist of? It would need to vary between a perfect score at one end, indicating that the network is completely sound, and a score of nil at the other, indicating total failure. In between, it should increase monotonically, in line with the overall health of the network.

If such a “Network Availability Index” could be regularly published in a trustworthy and neutral way, ILS prospectuses could be written with the value of the Index used as an indicator for catastrophic failure. If, for example, an Index covering the City of London was used in the writing of an ILS, the issuer might choose to define a 30% outage (however measured) as constituting a catastrophe, and therefore triggering the payment.

The Index should be available to whoever wishes to subscribe. Any subscriber could monitor it, in order not only measure catastrophic events, but also to use as a measure of general disruption. A 5% outage may not be enough to be defined as catastrophic, but the value may nonetheless be high enough for a subscriber to take some protective action of their own.

The mechanism would have to be trustworthy and independent. When investors buy into a natural-catastrophe ILS, they can be confident that the source of the underpinning meteorological data is entirely neutral. A cyber-catastrophe ILS would need to use a similarly reliable trigger.

B. Polling The Network

A computer network consists of a “mesh” of nodes. In a “mesh” structure, not every node is connected to every other node. Instead, nodes are linked by intermediaries. In the case of computer networks these “middlemen” are dedicated routing servers, which propagate data from node to node.


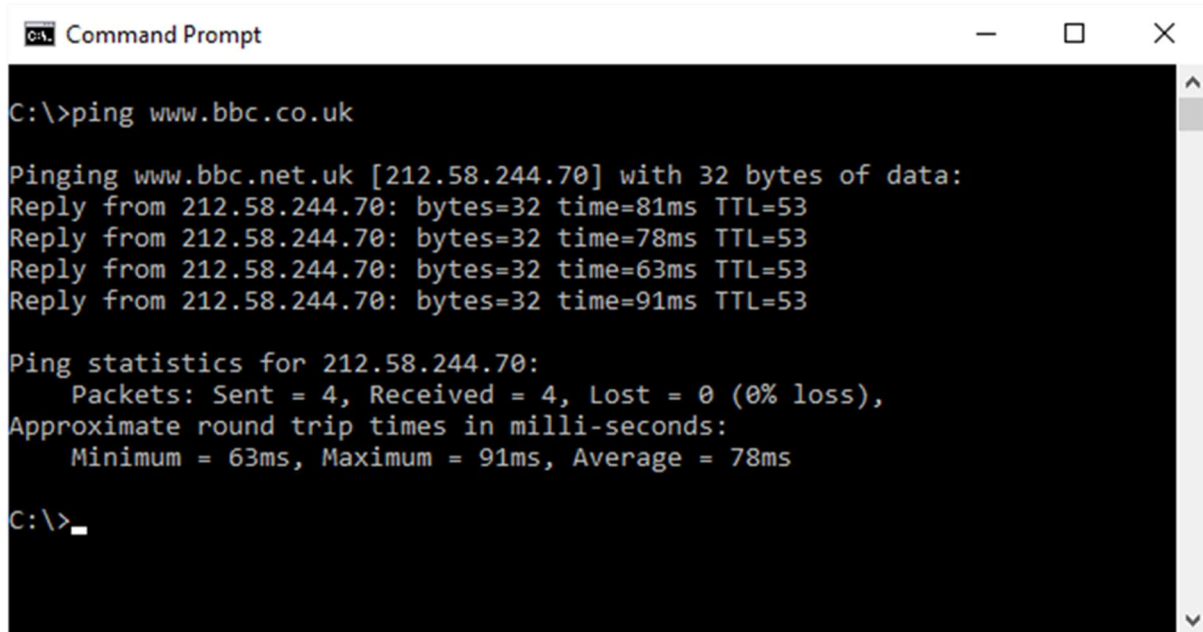


Figure 5 : A mesh-style Network

The simplest way of sending a message to a specified destination machine is to use a “ping”. This is a radar-like pulse – a small packet of information – which, if it can, will find the destination machine, and be returned to the point of origin as confirmation that the destination machine is operational and on the network. For example, pinging the BBC web server is a simple matter of opening a terminal and typing: `ping www.bbc.co.uk`. This sends four consecutive packets of data to the web server. If the server is up and running, then it returns a positive result as follows:



```
Command Prompt
C:\>ping www.bbc.co.uk

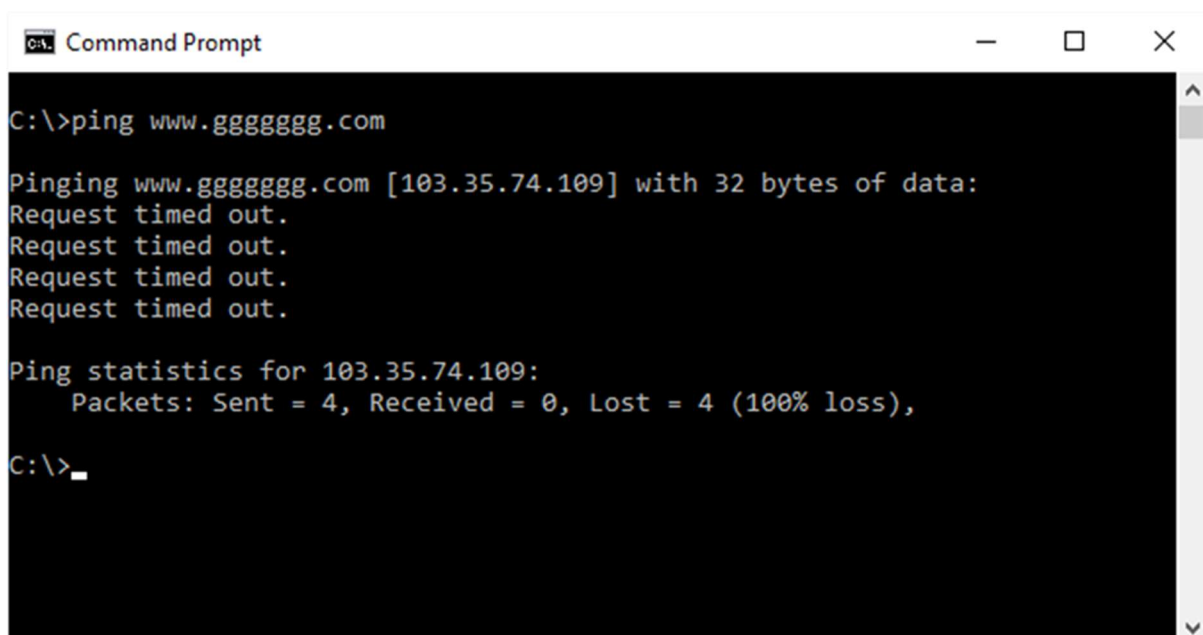
Pinging www.bbc.net.uk [212.58.244.70] with 32 bytes of data:
Reply from 212.58.244.70: bytes=32 time=81ms TTL=53
Reply from 212.58.244.70: bytes=32 time=78ms TTL=53
Reply from 212.58.244.70: bytes=32 time=63ms TTL=53
Reply from 212.58.244.70: bytes=32 time=91ms TTL=53

Ping statistics for 212.58.244.70:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 63ms, Maximum = 91ms, Average = 78ms

C:\>_
```

In order to contact the given web address, it must first be resolved to an IP address, by means of a Domain Name Server (DNS). This system is the global phonebook for all websites running on the planet.

If the website is found by the DNS, but is down, then the following is returned:



```
Command Prompt
C:\>ping www.gggggggg.com

Pinging www.gggggggg.com [103.35.74.109] with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 103.35.74.109:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\>_
```

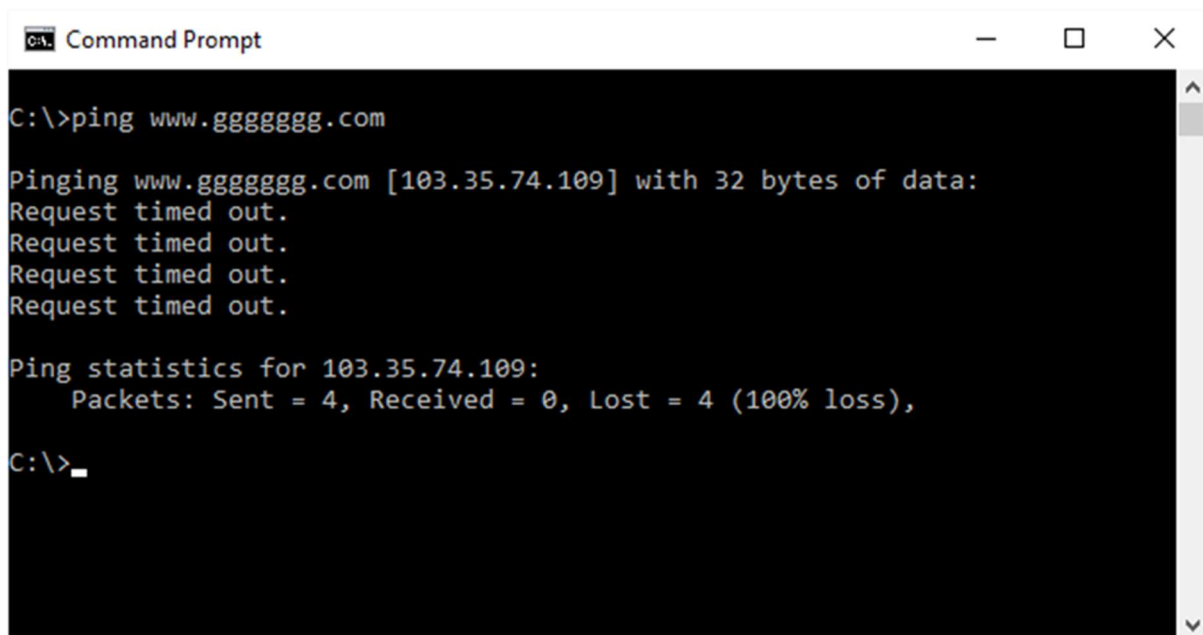
The DNS system is the means by which remote addresses are found every time a request is sent out over the internet.⁷⁴ In order to resolve the web address to an IP address, the protocol contacts DNS after DNS until it find the IP for the

⁷⁴ https://www.verisign.com/en_US/website-presence/online/how-dns-works/index.xhtml

domain in question. Once the IP address is known, an ICMP (Internet Control Message Protocol) Echo Request is sent out over the TCP/IP protocol, which is used by everything communicating over the internet.

Guiding the message to its destination are dedicated routers. The message proceeds in a series of hops, as each router forwards the message to the next appropriate router, using lookup tables.⁷⁵

If the host cannot be found by the DNS at all, the following is returned:

A screenshot of a Windows Command Prompt window. The title bar reads "Command Prompt". The command prompt shows the following text:

```
C:\>ping www.gggggggg.com

Pinging www.gggggggg.com [103.35.74.109] with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 103.35.74.109:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\>_
```

One way of measuring a network's overall availability would therefore be to repeatedly ping a representative sample of web addresses. The percentage of addresses which return the ping would then be an indicator of the overall network health.

However, it turns out that many web servers nowadays block ping requests. This is often done because of overly stringent security requirements. We will therefore use HTTP polling instead of ping, which works much the same way. How many addresses would need to be polled? Because of the mesh nature of the network, the path that the polling signal takes to its destination is not predictable. In general, there are a large number of possible paths between node A and node B. Not only that, the route from A to B is not guaranteed to be the mirror-image of the route from B to A. In order to assess the general health

⁷⁵ <https://www.metaswitch.com/knowledge-center/reference/what-is-ip-routing>

of the network, it would be wise to poll a large number of addresses in order to balance out any non-linearities.

Polling web addresses like this effectively conflates two tests: one of the DNS system and one of the routing system. The DNS system resolves the web address to an IP address, and then the routers find the node with that IP. If DNS servers go down, the test will fail, even if the network is still capable of propagating messages to and from IP addresses. This is acceptable – a widespread DNS failure, if large enough to cause catastrophic outage, certainly constitutes a Business Interruption cyber-catastrophe.

C. Constructing The Index

One way to construct and publish the Network Availability Index is as follows:

- Every few minutes, a central server polls a fixed list of domain names.
- Each domain name on the list either returns an echo within the timeout period or it doesn't.
- The central server calculates the proportion of all responsive domains.
- This proportion is the Network Availability Index. It is published on a website for anyone to see.

If the index is to be used in diagnosing cyber-catastrophe, it is important that the scope of the index is defined. One can imagine an index based on the health of the network backbone of the City of London. This would potentially be used to support an ILS which allowed for hedging of a buyer's exposure to the hardware relied upon by UK financial institutions. Therefore, the index should ideally be a measure of the availability of only the parts of the network used by those institutions.

Similarly, if an index consumer was interested in the portions of the network used by government, say, then polling different Ministries and Departments would accomplish this. An index based on the institutions belonging to the Critical National Infrastructure⁷⁶ would serve as an indicator of the health of the crucial sub-networks which allow communication between the most vital institutions.

However, computer networks do not operate in single locations. Thanks to fibre optic technology, the physical distances between pairs of nodes can be huge. A

⁷⁶ <https://www.cpni.gov.uk/critical-national-infrastructure-0>

single “hop” could cross the Atlantic. For the purposes of networking, the topology of the network is more relevant than the fixed geometry.

But a Network Availability Index needs to be localised to a specific geographical location if it is to be useful. Since a cyber-catastrophe is vanishingly unlikely to be spread worldwide, there is no point having a Network Availability Index without physical boundaries.

If we wanted to limit the Index constituents to servers within the City of London, as long as both the poller and the institutions polled were both within that geographical space, the local network would be used. The DNS servers and routers would be nearby (both topologically and geographically), and therefore would be part of the test.

Network routing algorithms tend to try the most direct route first. If the node being routed to was close to the source of the signal, the return of the polling signal would be an indicator of the existence of the local connection, which is what the Index is supposed to measure. The poll timeout would have to be set sufficiently low, so that the router would not have time to go and try more roundabout paths. The test must remain confined to the local network.

D. Constituent Weighting

The above scheme assumes that the index is calculated as a simple average of the availabilities of a large number of server addresses. One question to be addressed is that of the weighting of the index constituents. Are some web addresses more important than others in any way?

There are two factors which are at play here: one hardware-based and one social. The hardware factor is related to the topology of the network. Certain nodes may be better connected to the network than others. This will mean that more paths will lead in and out of those nodes. Therefore, when polling those nodes, the redundancy of the links will mean that a signal may reach that node even if half the network is down. Therefore, the usefulness of such a node in determining the overall availability of the network is (perhaps counterintuitively) diminished.

Relatedly, however those very nodes will be better connected on account of their having many more users than other, less well-connected nodes. Therefore,

the social cost of such an outage for that area of the network is larger than for other, less interconnected areas.

When measuring Network Availability these two factors could be said to balance each other out. After all, the connectedness of a certain address to the network is put in place as a function of the number of users who connect to it. In order for that address to perform at a typical speed, the number of notional users per notional connection must be similar to the number of users per connection across the network as a whole. Therefore, provided that a sufficiently large number of addresses over a sufficiently large network area are included in the Index, there is no need to adjust for high connectivity or heavy use.

E. Proposed Architecture

With all the above in mind a simple scheme might look like the one illustrated in Figure 6. A central server polls the network, the index is centrally calculated and published as a number to (say) a website.

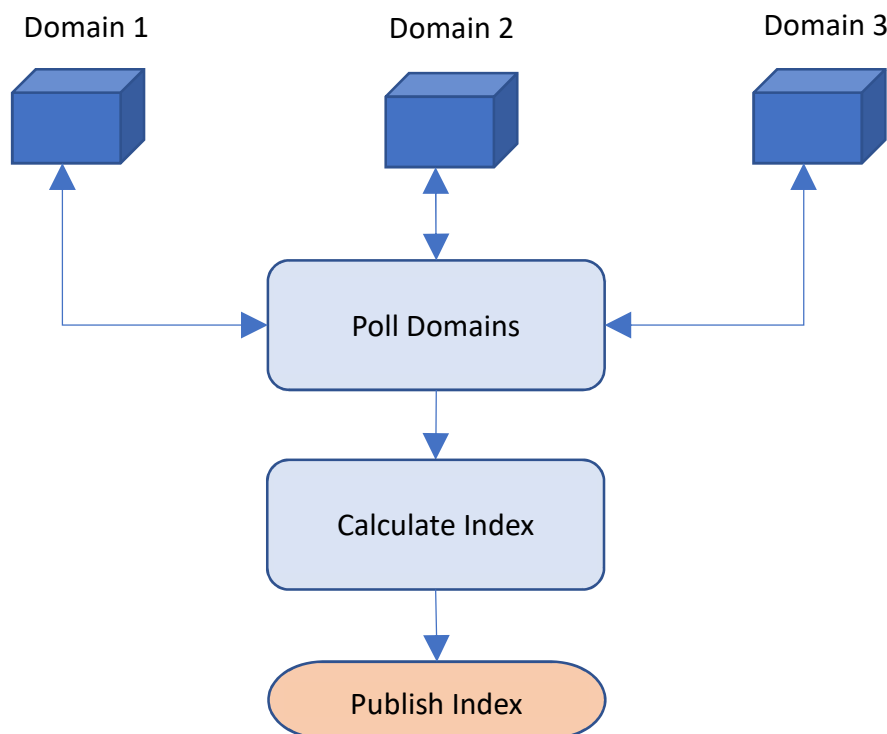


Figure 6 A simple index publication setup

One drawback to this scheme is that it is susceptible to the very network outages it is measuring. If the paths to the central server are broken, the index will not be visible to subscribers.

This can be remedied with the use of a distributed system. Here, the central server publishes the index to a number of “receivers”, who then can publish their results to their own dedicated websites. So if the user cannot locate receiver, others are available.

This scheme is illustrated in Figure 7.

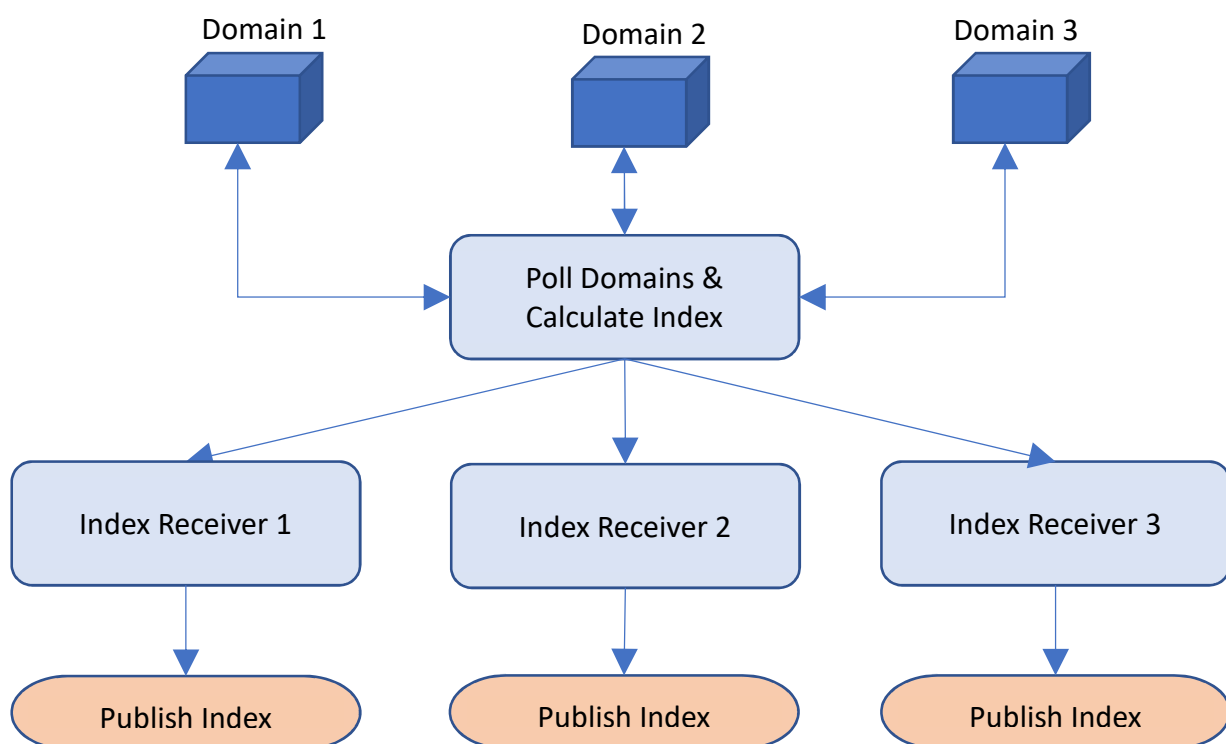


Figure 7 Index publication using receivers

However, the same concerns which affect the receivers also affect the central polling machine. If the poller is geographically close to the servers being polled, there is an increased risk that, in the event of a network outage, the poller will be unable to publish the Index at all. A way of mitigating this risk is to use a system of distributed pollers, as shown in Figure 8.

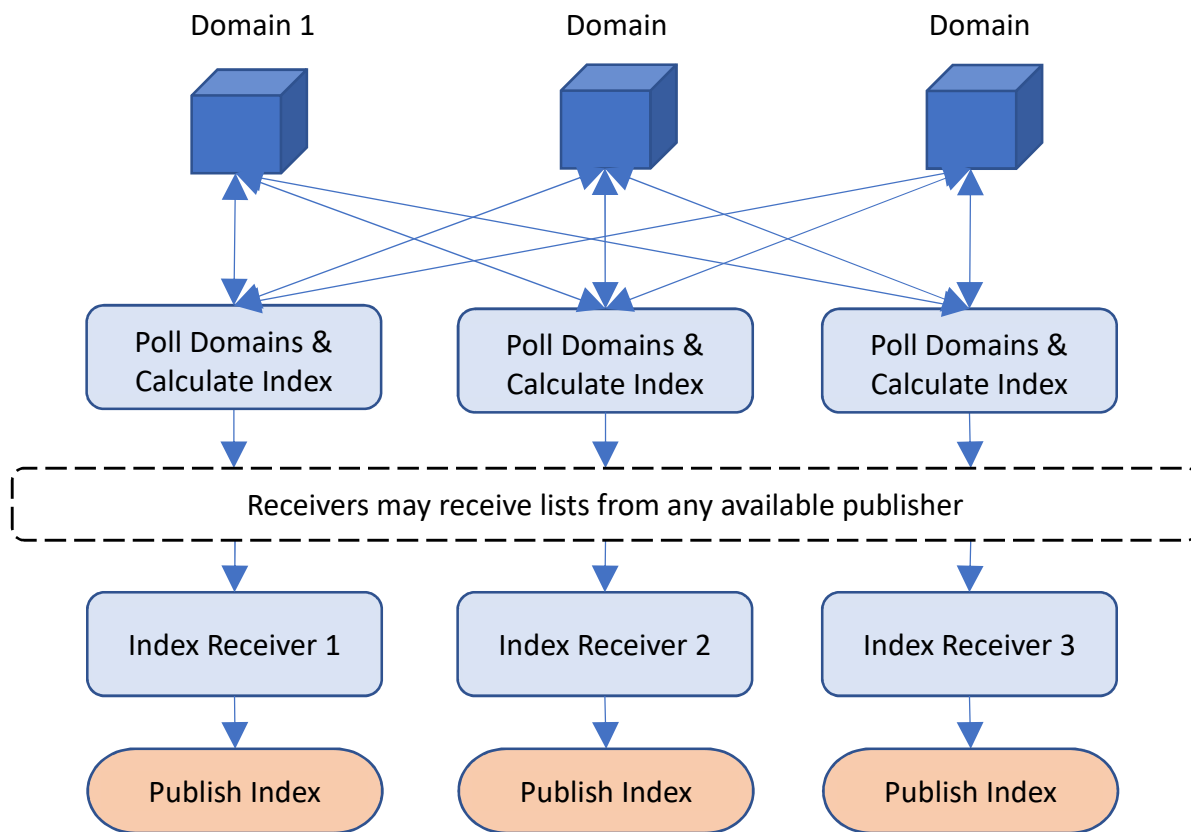


Figure 8 Publishing using multiple polling servers

The next drawback to be addressed is that the index available to every user is the same index. If different users require different information, there is no possibility of them either drilling down to the constituent data, or requesting different versions of the index.

In order for this to be possible, each receiver needs to have access to the constituent data. However else the system works, the index needs to be constructed at the receiver level, as illustrated in Figure 9.

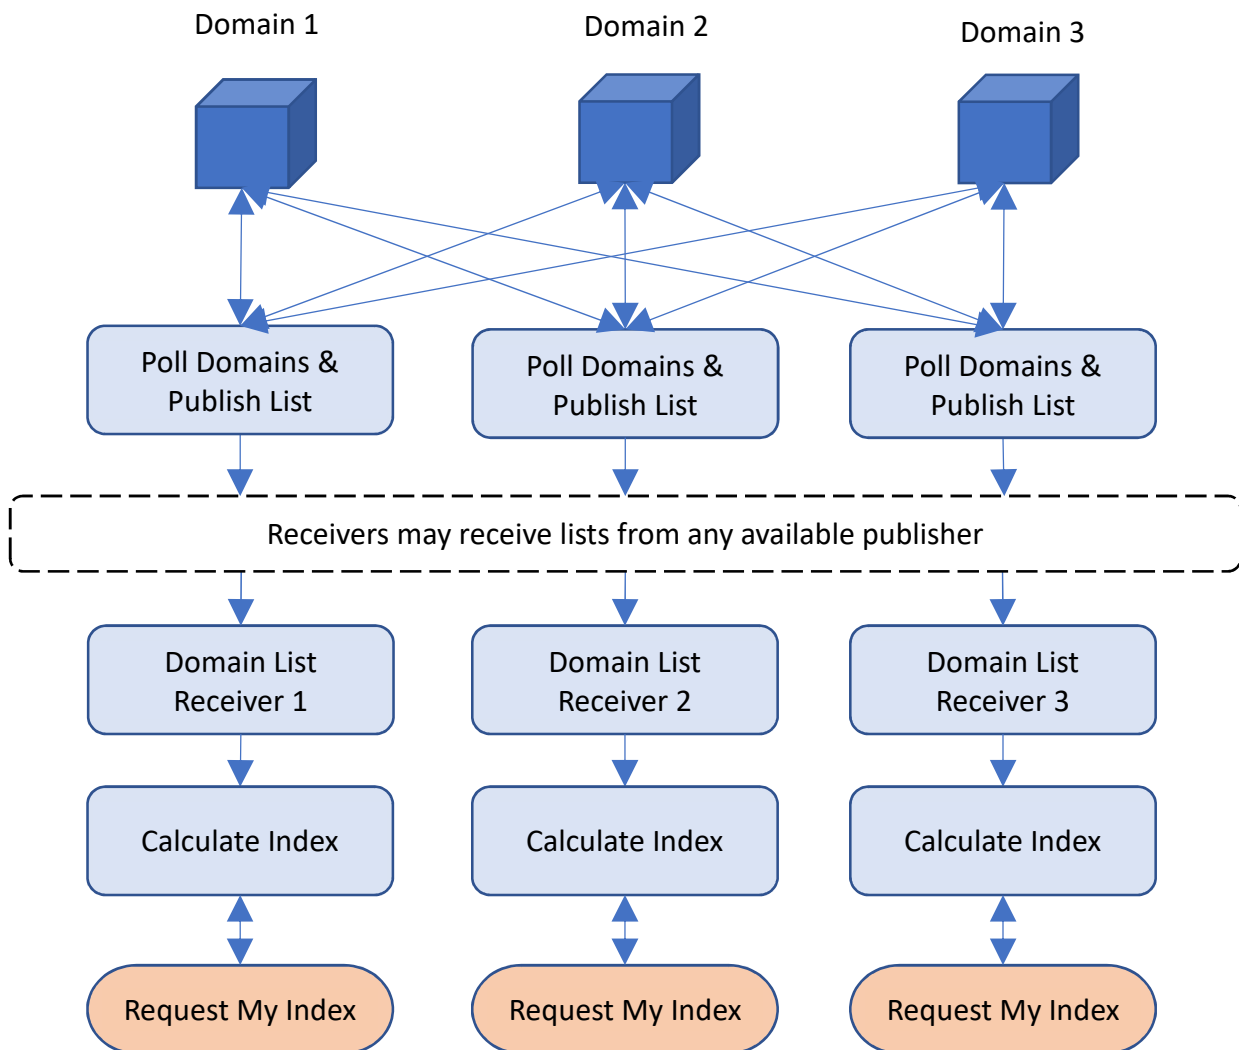


Figure 9 Receivers receive all the list of results and calculate the index on demand

F. A Broadcast/Receive Solution

What is required, therefore, is:

- the ability for a single machine, or a distributed set of machines, to poll a set of addresses in nearby geographical location, and
- publish the poll results so that receivers can subscribe to one or all of the polls, and
- for each receiver to be set up to calculate the index itself, and
- for bespoke indices to be able to be calculated on a per-user basis.

Broadcasting The Poll Results

An ideal technological model for the broadcast portion of the architecture is known as “broadcast/receive”.⁷⁷

In the simplest form of such a system there is a single machine called a transmitter, and zero or more receivers. If a user wishes to broadcast, they send a message request to the transmitter (via HTTP, say), which then broadcasts the message to any receivers which happen to be listening at the time. The transmitter does not know how many receivers (if any) are listening, nor does it get any kind of receipt of the message. The message is simply sent out to whoever is there.

One implementation of this paradigm is the ChainZy model.⁷⁸

As shown in Figure 5, the risk of network outage necessitates the existence of not only multiple receivers, but also multiple transmitters. This facility is also provided by ChainZy in the form of their “woven broadcasting” scheme. Under this scheme, multiple transmitters are run and therefore multiple ledgers are created. However, the “woven broadcasting” method allows the data to be stitched together.⁷⁹

Fortunately, the receivers in our scheme do not necessarily have to do anything difficult to combine the different sources of published data. Each receiver could subscribe to all publishers they know about, and combine the published lists into a master list. If an address was reachable by at least one of the pollers, it would be registered as a “pass”. This renders the system a good deal more robust. An address must be unreachable from multiple different nodes in order to be marked as a “fail”.

Monitoring The Index

As we have sketched, users need to be able to request a calculated index relevant to them from each receiver.

⁷⁷ Otherwise known as “publish/subscribe” or “pub/sub”. See

https://en.wikipedia.org/wiki/Publish%E2%80%93subscribe_pattern

⁷⁸ <https://www.zyen.com/work/types-work/mutual-distributed-ledger-aka-blockchain-technology/chainzy/>

⁷⁹ <http://www.zyen.com/what-we-do/1615-chainzy.html>

Each receiver, therefore, will have on board some “index construction” code. These programs will be long-running services, which wait for user requests. Each request will contain the identity of the user. On receiving a request, the Index Constructor will look up the user’s associated formula (more on this below). The calculated Network Availability Index is then returned to the user.

Why would we want different users to have access to different forms of the index? There are a number of possible reasons:

- We may set up the list of polling addresses to cover a large geographical area. The user may only be interested in the index which is constituted from subset of polled addresses.
- Similarly certain users may be more interested in addresses which suffer from a cloud outage. Their requested index might therefore cover the clients of a specific cloud provider.
- We may wish to privilege some types of users over others. Those who have a direct interest in the status of the index-based ILS – i.e., the reinsurers and investors – may pay more for information which is made available sooner or in a more complete way.
- Potentially, the index should not be made available “live” to any entity who owns one of the addresses polled. If it is public knowledge that a Network Availability Index is reaching the trigger point for an ILS, the owner of an ILS could shut down their own address in order to tip the index over the crucial point. This element of moral hazard could also be addressed by simply ensuring that the pool of polled addresses is large enough.
- Other parties with an interest in catastrophe are governmental or crime-detection agencies. These institutions could therefore receive the full Index in a privileged or more timely way. Possibly they could run their own receivers so that they don’t need to request information, but simply subscribe to it.

In order for each user to have the ability to request its own version of the index, the Index Construction code on board the receivers need to be able to look up a specific formula for each user. The broadcasters therefore need a separate channel available, in order to push daily updates to their receivers about the user universe, each user’s cryptographic key, and the index formula relevant to each user. To optimise performance, this information can then be stored on each receiver, ready for users’ requests.

Triggering The Payment

The key users of the system would be the reinsurers, as issuers of Cat Bonds, and investors in those Cat Bonds monitoring their triggers. A reinsurer would therefore subscribe to their particular version of the index, and when the index breaches a certain level as specified in the ILS definition, ensure a quick payment to the insurers to cover catastrophe payments.

One of the most exciting uses of ledgers is in the use of smart contracts – or code which runs automatically given a certain set of condition. These programs sit on a ledger, fully available for inspection by interested parties.

This technology provides a way of automating the process from end to end. The ILS could exist as a piece of code on a ledger somewhere. The code would be the “user” which requests the Network Availability Index from the Index Calculation servers. Once the index is received, the code on the ledger would compare it to the triggering level defined in the ILS specification. If the level has been breached, the code would automatically perform a payment to the insurers who need it. This payment could be in the form of a cryptocurrency stored on the ledger, or an ordinary electronic payment.

Index Updates

It is necessary for an index like this to remain current. If a web domain goes down, not for network failure reasons but because of a server-side crash, or because the site is now no longer maintained, the index will be affected. If a site is consistently down over, say, a 24-hour-period, then an alert can be raised in order for the index owners to perform a manual check. If the domain is indeed down, then it is no longer a good indicator of network performance and should be dropped from the index. If some user has set up a trigger, and the network availability is close to the trigger level, then it is important that any bad domains are removed from the index so that they do not trigger a payment.

G. The Benefits Of This Approach

The ChainZy technology is a form of Mutual Distributed Ledger (MDL), with some differences. The key difference is that the technology makes a distinction between broadcasters and receivers. In a typical distributed ledger situation, there is no difference, and all participating nodes are equal. If a user wants to add data to the ledger, they first add it to their local copy. The data is then

propagated to other nodes. Data is only added permanently to the chain when there is a consensus among nodes that the data added was valid.

However, if we wish to build up an authoritative published index, a standard MDL is not suitable. Some level of centralisation, as with ChainZy, is required in order to ensure that polling machines sit at appropriate points in the network. This also makes the system much simpler and easier for users to understand.

In addition, because we will be able to specify the number of receivers and their location in the network, we will not have to simply hope that we have enough users running nodes at appropriate points. ChainZy's broadcast/receive model is therefore ideal for this application.

So far, we have limited our discussion to the concept of cyber-catastrophe. As mentioned in previous chapters, currently ILSs are mostly used to offload risk of natural catastrophe. How might the technology we have discussed be used in other applications?

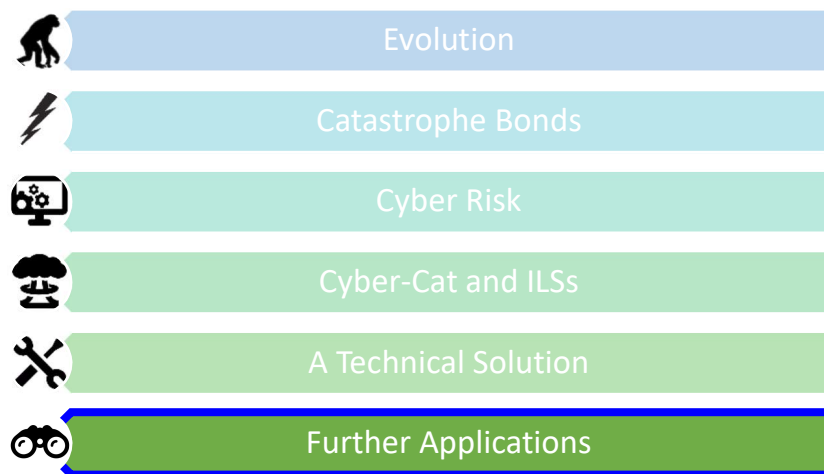
6. Further Applications

In the last few chapters we have described how to build an index based on distributed data collected by a distributed system. Thanks to our use of Smart Ledgers, the index is publicly auditable and cannot be influenced by malicious parties.

In the architecture outlined in the last chapter, there are three separate stages:

1. The use of polling to build up a picture of a system's general availability, expressed as a percentage,
2. The ability for subscribers to request "tailored" aggregations of the data, based on their individual permissions and interests, and
3. A way to use an index's published value to trigger a transaction or an action in general.

In this chapter we will look at all of these sub-architectures in turn, in order to see where else they might be applied.



A. Polling

In previous chapters we have described a polling system which operates over the whole of a network, treating the addresses to be polled as equal to one another in importance. We have also treated each node as having a binary state – it is either responsive or it isn't. The polling algorithm can be made more inventive in order to gain more insight into the availability of the network.

Network Speed

One aspect of the network worth drilling down into is its speed. Any decent polling system can provide data on how long a given address took to respond to a data request. The length of time a signal takes to return could be used as an inverse weighting on the index. This would mean that the index took into account effects caused by inefficient routing or heavy traffic when calculating a measure of network availability.

Organisational Reporting

In what situations might an organisation find this scheme useful *within* that organisation?

Any large enough network is susceptible to chaotic/systemic failures which cannot be predicted. Multiple small errors can occasionally conspire to make a “perfect storm” effect. A distributed pub/sub system can measure the health of a system in a holistic way, rather than trying to home in on specific failures. The affected portions of the system can then be isolated stage by stage.

Moreover, there exist organisations whose cyber resilience is of national importance. If an organisation is required to publish statistics of fully auditable network availability on a regular basis, then this could provide a useful data point for government and other interested parties. This would be a more useful measure of technological resilience than obfuscatory reports on IT spend. Large organisations such as the NHS, the electricity infrastructure and the military could be candidates for this.

Mutual Polling

As we saw in Chapter 1, before the Industrial Revolution, the provision of fire insurance was dominated by mutuals, who had the ability to assess each of their members.

Something similar could be done through the use of polling. Without involving external insurers, a set of interested parties could set up a distributed polling network, polling each other’s domains and other domains in which they shared an interest. All parties would contribute to the pool and cover each other’s BI losses.

As with fire insurance, it would be important to ensure that the risk borne by each member was sufficiently decorrelated. Mutuels declined among close neighbours since fire risk in big cities increased, and was highly correlated between members of the same neighbourhood. If, for example, half of the mutual's members used the same cloud provider, then the concentration of risk might be deemed too extreme.

B. Publishing Indices

Tailoring

In the system we described in the previous chapter, the receivers submit their identities and an index is built for them. This is a concept which could be applied to any index or aggregation, not just those linked to insurance.

By using Smart Ledgers, an index can be tailored to each individual's requirements and/or permissions. Smart Ledgers allow the publication of user-specific indices, with user-specific timing lags, without having to give any users access to the underlying data. The equations used to create the index are visible but only the single number is published.

Provided that the input data was provided in a consistent way, the Index Construction code running on the receivers could calculate and publish the index in any way a user wanted. This creates an opportunity for a business model where a customer does not have to buy raw data and perform calculations, but gets personalised metrics purchased one by one.

Input to the system would need to be standardised. For example, if the ILS in question is based on natural catastrophe, the weather data would need to be sourced from a recognised authority in a universal format. If the trigger for the ILS was health information (in the wake of some pandemic, for example), then the news of fatalities and refugee movements would need to be validated and authoritative.

This standardisation is becoming easier and easier. The last few years has seen boom in the number of programmatic APIs published by weather and news bureaus, social media feeds, and financial markets reporting. These can be subscribed to by the receivers running the Index Calculation code, in order to publish bespoke indices to the users

Data Sharing

Tailored indices are ideal in any situation where there are a number of actors who have different permissions for different views of the data. Data privacy is a challenge in the regulatory world. The regulators want to be able to assure the institutions that they are auditing that their information is confidential. At the same time, the regulators may need to share the data with other bodies, and publish suitably aggregated versions of it.

Recently Brazil's central bank tackled this problem by using a blockchain system.⁸⁰ This is a useful way of sharing data, but cannot provide different levels of access to different users in an intelligent way. Using Smart Ledgers would be a powerful way to give each party the correct slice of the data, whether that be access to certain raw data points or only to aggregate numbers.

Data Anonymisation

A key aspect of the model outlined above is the ability to publish bespoke, partial information. In the world of medical research, when data is provided to researchers, it very often needs to be anonymised and aggregated beforehand, in order that confidentiality is not breached. However, it may be different research requires different pre-processing, either because different aggregations are required or because the requester has different permissions. It is tedious to do the aggregation specially for each requester. The ability for the publisher to publish a tailor-made aggregation on hidden information would allow different receivers to request the aggregations they were entitled to.

Media

Advertisers face new challenges thanks to the transformation of electronic media. First, they need to target their ads in a more and more fine-grained way; only displaying each ad to someone who is most likely to respond based on their profile or demographic. Second, when a user is engaged with a platform it needs to respond quickly and dynamically, and present the ads before the user moves on.

Media networks therefore should be able to provide data to advertisers about their consumers in real time. However, the data is valuable and confidential,

⁸⁰ <https://www.coindesk.com/brazils-central-bank-plans-blockchain-data-exchange-for-regulators/>

and will not be made available to advertisers in raw form. Instead, metrics encoded on Smart Ledgers will process the data and publish the resultant figures to the advertiser.

The system provides a remarkable bonus here. The Smart Ledger code can be specified by the receiver. An advertiser could therefore load their quantitative models on to the Smart Ledger, which would run on data that they would never see. This squares the circle of not allowing customers to the base data, but allowing them to run their own processes on it.

Sub-Network Data

We have assumed so far that the subscriber doesn't have to know too much about the polling mechanism which lies behind the index calculation. It may be, however, that they want to leverage the power of the polling system to gain further insights.

A customer could be interested in the performance of certain areas of a network – say, those nearest a certain Domain Name Server – and attach more or less importance to those areas. Their bespoke index would then correspond to specific positions in the network topology.

This would allow a customer with a special interest to analyse parts of a large network and look for correspondences between certain areas of failure. If the failure of a certain bank is always associated with the failure of, say, a certain embassy, this may indicate some kind of political risk.

Similarly we can assess the health of the network divided into sub-networks. The network can be cut up in any way – geographically, topologically, by industry. The ability to find correlations between the health of disparate sectors is something which the financial markets would gladly pay for. If historic data indicates that a cloud outage affecting a certain retail chain is associated with outages affecting other areas, then a financial institution might choose to hedge this risk ahead of time.

C. Triggers

Weather Derivatives

While ILSs are currently used to trigger payments in the cases of natural catastrophes, there is no reason why they should not be used in simpler weather-based financial instruments. There is an active market in “weather options”. These can be used as hedges in the energy industry, ensuring that an unusually warm winter does not result in large losses. The retail and travel industries can also suffer in the case of adverse weather, and these losses can also be offset. Practically, there is no difference between hedging the weather and hedging against an unfavourable swing in an exchange rate.

Weather derivatives use indices published by centrally agreed institutions such as Earth Satellite Corp. Hooking this data up to Smart Ledger technology would provide a means of ensuring automatic, timely payments to option holders.

Other Applications

Such an argument is not confined to weather derivatives, of course. All financial derivatives contracts are conditional payments based on a centrally agreed reference value. Smart Ledgers could be used in any similar situation. The question repeatedly asked about the technology is whether some other existing technology can't do it just as well. Admittedly, then, for simple derivatives, there is probably not much to be gained through the use of Smart Ledgers.

However, when an event is conditional upon centrally agreed indices, the market is reliant upon those indices being available to be referenced. Complexity is introduced when a derivative is triggered by a bespoke, tailored index. In this case, Smart Ledgers provide a way for market participants to design their own index and set up their own triggers, without needing to own the underlying data.

Conclusion

In everyday insurance claims, the insurer looks to clarify the causes and effects, before releasing any payments. This is because it is vital that the insurer makes certain that the claimant is not trying to defraud them.

However, in the case of natural catastrophes, while the loss adjuster must establish the correct levels of compensation, the causes are already established. And, through the use of independently verified data, the re-insurer can be quick to pay their customers.

A similar attitude may have to be brought to bear on cyber-catastrophe risk. The complexity of the system often means that the cause of a massive cyber failure takes months or years to discover. Nevertheless, the industry cannot waste time attempting to discover the cause of the problem before payment is made.

It is also vital that the insurance industry unpicks the notion of “cyber” risk, clarifying for all parties precisely which risks are covered and which are excluded. Once this is done, they will be in a position to evaluate and offload their risk of cyber-catastrophe, in the form of Insurance-Linked Securities, to the capital markets.

The combination of adopting a quick, programmatic attitude to payment, accurately parametrising the risk, and using of capital markets makes the setup ideally suited to Smart Ledger technology. The Smart Ledger approach allows participants to measure the current levels of network risk, and potentially trigger insurance payments. The benefit of this technology is that it provides high security and large degrees of flexibility in structure, yet prevents the rise of an over-weening central third party, thus encouraging competition and innovation in the provision of cyber insurance direct to clients.

The distributed architecture of a Smart Ledger system is robust enough to withstand all but the most severe cyber-catastrophes.

Principal Authors



Sam Carter

Carter Research Ltd.

A financial services researcher and programmer, with a special interest in quantitative development and natural language processing. He has worked for two decades in the City of London as a developer, quantitative analyst and product manager, in the capital markets and fintech industries. During that time, he has managed and coded the development of systems for reference data, bond and portfolio metrics, synthetic collateralised debt obligations (CDOs) risk calculations, commodities pricing, and automated compliance. He holds an MSc in Financial Mathematics from King's College London, where he worked on credit default swaps, and an MA in Linguistics from UCL, where he worked on information theory in the syntax of natural language. He is interested in the interfaces between the different actors in the financial world – the dialects spoken by people in finance, mathematics, legal, compliance and trading – and what information is lost when those worlds attempt to communicate with each other. Sam is an amateur guitarist and pianist. He speaks terrible French and even worse Afrikaans.



Professor Michael Mainelli FCCA FCSI FBSC

Executive Chairman, Z/Yen Group

A qualified accountant, securities professional, computer specialist, and management consultant, educated at Harvard University and Trinity College Dublin, Michael gained his PhD at London School of Economics where he was also a Visiting Professor. He began his career as a research scientist, later becoming an accountancy-firm partner and a director of Ministry of Defence research. During a spell in merchant banking in 1994, he co-founded Z/Yen, the City of London's leading commercial think-tank. He has led Z/Yen from creating Smart Ledgers (aka blockchains) through the Financial Laboratory, Taskforce 2000, Long Finance, Global Financial Centres Index, Global Green Finance Index, and Global Intellectual Property Index. He is a non-executive director of two listed firms and a regulator, Emeritus Professor at Gresham College, Fellow of Goodenough College, and a past Master of the Worshipful Company of World Traders. His third book, **The Price of Fish: A New Approach to Wicked**

Economics and Better Decisions, co-authored with Ian Harris, won the 2012 Independent Publisher Book Awards Finance, Investment & Economics Gold Prize.

Acknowledgments

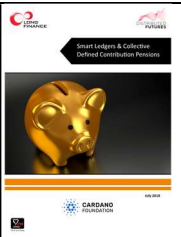




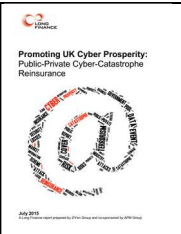
We would like to thank Julian Enoizi, Ian Coulman and Steve Coates from Pool Re for their patient help in the area of cyber-terrorism and insurance.

Mike St John-Green provided fascinating information on the topic of cyber-security.

Mark Duff provided vital background on the insurance industry, and pointed to previous work done on the cyber problem.

And Ben Morris of Z/Yen realised the initially sketchy demo ideas with surprising ease.

Other Long Finance Publications

Title	Authors	Year	Publisher
	Iain Clacher, Con Keating, and David McKee	2018	Long Finance (July 2018), 47 pages.
	Sam Carter	2018	Long Finance (June 2018), 55 pages.
	Centre for Economics and Business Research	2018	The Worshipful Company of World Traders and Long Finance (April 2018) 78 pages.
	Professor Tim Connell and Bob McDowall	2018	Long Finance (March 2018), 102 pages.
	Maury Shenk	2018	Long Finance (February 2018), 61 pages.
	Michael Mainelli, Chiara von Guntzen, and Mark Duff	2015	Z/Yen Group, Long Finance (July 2015), 50 pages.



Distributed Futures is a significant part of the Long Finance research programme managed by Z/Yen Group. The programme includes a wide variety of activities ranging from developing new technologies, proofs-of-concept demonstrators and pilots, through research papers and commissioned reports, events, seminars, lectures and online fora.

Distributed Futures topics include the social, technical, economic, and political implications of smart ledgers, such as identity, trade, artificial intelligence, cryptography, digital money, provenance, FinTech, RegTech, and the internet-of-things.

www.distributedfutures.net



Cardano Foundation is a smart ledger and cryptocurrency organisation based in Zug, Switzerland. The Foundation is dedicated to act as an objective, supervisory and educational body for the Cardano Protocol and its associated ecosystem and serve the Cardano community by creating an environment where advocates can aggregate and collaborate.

The Foundation aims to influence and progress the emerging commercial and legislative landscape for blockchain technology and cryptocurrencies. Its strategy is to pro-actively approach government and regulatory bodies and to form strategic partnerships with businesses, enterprises and other open-source projects. The Foundation's mission is the promotion of developments of new technologies and applications, especially in the field of new open and decentralised software architectures.

www.cardanofoundation.org



“When would we know our financial system is working?” is the question underlying Long Finance's goal to improve society's understanding and use of finance over the long term. Long Finance aims to:

- ◆ expand frontiers - developing methodologies to solve financial system problems;
- ◆ change systems - provide evidence-based examples of how financing methods work and don't work;
- ◆ deliver services - including conferences and training using collaborative tools;
- ◆ build communities - through meetings, networking and events.

www.longfinance.net



Z/Yen is the City of London's leading commercial think-tank, founded to promote societal advance through better finance and technology. Z/Yen 'asks, solves, and acts' on strategy, finance, systems, marketing and intelligence projects in a wide variety of fields. Z/Yen manages the Long Finance initiative.

Z/Yen Group Limited
41 Lothbury, London EC2R 7HG, United Kingdom
+44 (0) 207-562-9562 (telephone)
hub@zyen.com (email)
www.zyen.com